ORNL/LTR-2022/399

The Second Artificial Intelligence for Robust Engineering and Science Workshop (AIRES 2)

Workshop Report

David E. Womble, Christy L. Hembree, eds

March 2022

OAK RIDGE National Laboratory

ORNL IS MANAGED BY UT-BATTELLE LLC FOR THE US DEPARTMENT OF ENERGY

DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via OSTI.GOV.

Website www.osti.gov

Reports produced before January 1, 1996, may be purchased by members of the public from the following source:

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone 703-605-6000 (1-800-553-6847)
TDD 703-487-4639
Fax 703-605-6900
E-mail info@ntis.gov
Website http://classic.ntis.gov/

Reports are available to US Department of Energy (DOE) employees, DOE contractors, Energy Technology Data Exchange representatives, and International Nuclear Information System representatives from the following source:

Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831
Telephone 865-576-8401
Fax 865-576-5728
E-mail reports@osti.gov
Website https://www.osti.gov/

ORNL/LTR-2022/399

Computing and Computational Sciences Directorate

THE SECOND ARTIFICIAL INTELLIGENCE FOR ROBUST ENGINEERING AND SCIENCE WORKSHOP (AIRES 2)

WORKSHOP REPORT

David E. Womble

and

Christy L. Hembree

editors

March 2022

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831
managed by
UT-BATTELLE LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

ABBREVIATIONS

AD	anomaly detection
ADCME	Automatic Differentiation Library for Computational and Mathematical Engineering
AIRES	Artificial Intelligence for Robust Engineering and Science
AM	additive manufacturing
APS	American Physical Society
ARPA-E	Advanced Research Projects Agency–Energy
ART	Adversarial Robustness Toolbox
ASME	American Society of Mechanical Engineers
AST	adaptive stress testing
CFD	computational fluid dynamics
CT	computed tomography
DL	deep learning
DNN	deep neural network
DNS	direction numerical simulations
DOE	US Department of Energy
DT	digital twin
Fermilab	Fermi National Accelerator Laboratory
FPGA	field-programmable gate array
GAN	generative adversarial network
GMPS	gradient magnet power supply
GPU	graphics processing unit
HPC	high-performance computing
IMMS	IoT-enabled maintenance management system
LANL	Los Alamos National Laboratory
LIME	Local Interpretable Model-Agnostic Explanations
LSTM	long short-term memory
MATI	Maintenance Advanced Technology Initiative
MIT	Massachusetts Institute of Technology
ML	machine learning
MNIST	Modified National Institute of Standards and Technology
MPI	Message Passing Interface
NRC	Nuclear Regulatory Commission
NLP	natural language processing
ODE	ordinary differential equations
ORNL	Oak Ridge National Laboratory
PDE	partial differential equation
physML	physics-informed machine learning
PID	proportional–integral–derivative controller
PINN	physics-informed neural network
PLM	product lifecycle management
PNNL	Pacific Northwest National Laboratory
RL	reinforcement learning
RNN	recurrent neural network
SADL-AT	Semantic Application Design Language Assurance Toolkit
Sandia	Sandia National Laboratories
SINDy	sparse identification of nonlinear dynamics
SISL	Stanford Intelligent Systems Laboratory
SNS	Spallation Neutron Source

UB	State University of New York at Buffalo
UCF	University of Central Florida
UQ	uncertainty quantification
UTK	University of Tennessee, Knoxville
V&V	validation and verification

1. INTRODUCTION

On January 19–21, 2021, 173 researchers convened virtually to discuss Artificial Intelligence in Robust Engineering and Science (AIRES). This was the second meeting in the AIRES workshop series and focused on the digital twin (DT).

The concept of the DT has become pervasive in engineering since the introduction of the term at the beginning of the current millennium. Unsurprisingly, the definition of DT is almost as varied as the number of applications put forward in the literature. The workshop did not attempt to reach a consensus definition but instead accepted a relatively generic definition of a DT, which serves as a digital representation of an engineered or natural system. This broad definition assumes that the objective of a DT is to facilitate better decisions—such as enhanced control of the subject system, prognostics, and maintenance decisions for engineered systems—or to respond to the current or predicted state of a physical system.

The DT lifecycle can be divided into three phases: design, construction or manufacturing, and operation. Depending on the application, these phases have subjective interpretations and different levels of importance. An important component of the engineering and science in each of these phases is the ability to develop an accurate and computationally tractable system model. For the design phase, the focus tends to be on developing and using first-principles models. Machine learning (ML) can be used to develop computationally efficient models to span multiple lengths and time scales (e.g., to close turbulence models or to develop constitutive models). The construction and manufacturing phase begins with the design and adds the inherent uncertainties and imperfections into the model. In turn, the model provides the initial input for the operational phase, and the challenge transitions to reliably collecting operational data and updating the models. In many cases, this is inherently an AI and ML problem because the model must be updated using operational data without the ability to inspect the physical system. This workshop focuses on the AI research challenges.

Again, the DT lifecycle (i.e., design, construction, and operation) must be implemented in a given subject domain. As with any system, periodic maintenance and enhancement of the DT ensures its longevity and the consistent value addition that it can provide.

2. WORKSHOP ORGANIZATION

The AIRES 2 workshop included a keynote speaker, eight other invited speakers, and several contributed talks. The workshop included attendees with expertise in the foundational aspects of AI and ML as well as in the application of AI in a range of disciplines. The talks were organized primarily around three themes:

1. construction of DTs,
2. application and deployment of DTs, and
3. techniques to provide assurance.

The workshop included twelve breakout sessions. Topics were solicited in advance from registered attendees, and sessions with sufficient interest—based on an online poll—were included in the workshop. Thus, the selection of the breakout topics themselves reflected the research priorities of the participants. Workshop participants also had the chance to nominate breakout leads.

The breakout session leads were asked to define and address a series of questions for their area and to produce a short report. Each report aimed to

1. record the questions that were identified to guide the breakout discussion,
2. identify the key challenges that must be addressed in the area,
3. summarize the state of the art,
4. identify specific research challenges, and
5. summarize the impact of successfully addressed research opportunities.

The breakout session reports are included in APPENDIX A.

The complete AIRES 2 program, including a list of speakers, their biographies, and the titles of their presentations is provided in APPENDIX B. A list of the attendees is included in APPENDIX C.

3. RESEARCH NEEDS FROM AIRES 2

Although the breakout sessions addressed different aspects of DTs, several consistent themes in the research challenges and priorities for DTs emerged and are described below. The six themes below link naturally with each other. For example, model construction and continuous learning are naturally intertwined, and anomaly detection can never be independent of assurance and robustness requirements. The six research challenge themes that emerged in the AIRES 2 workshop are:

1. **Assurance**. The most referenced research need identified in the breakout sessions was *assurance*. Assurance for robust and reliable DTs is important in all applications, and the need is not limited to safety-critical systems. At an intuitive level, assurance addresses the question of whether an AI is making the right decision for the right reason. Assurance is a broad area that includes everything from uncertainty quantification (UQ) to causal inference. Within the assurance stack, four areas were identified as the highest priorities by workshop attendees:

    a. UQ. Uncertainty is an intrinsic part of both the data and the model, and rigorous bounds must be computed to guarantee a robust and reliable DT. It was noted that uncertainties that arise in the training phase become model uncertainties during the inference phase.

    b. Validation. Validation of a DT considers the appropriateness of the model and can only be considered in the context of the intended application. Validation must be a continuous process with the evolving state of the physical system. Any validation process must consider the appropriateness of training and inference data, the form of the AI, and the training process.

    c. Robustness. Robustness often refers to how the ML model responds to small changes in the data. For example, if the training data is relatively close to the operational data, then the DT should return results that are close. Robustness for the DT and AI more generally depends as much on the selection of the data and measures of closeness as it does on the model and training. For the DT, this must be expanded to include the full workflow, robustness to adversarial attacks, and unexpected occurrences in the environment or the data.

    d. Explainability and causal analysis. Explainability focuses on the human-computer interface and the ability of an AI or DT to explicitly associate a decision with a specific meaningful correlation identified in the data. Causal analysis goes significantly further and attempts to identify the causal relationships that underlie the identified correlations. Establishing these causal relationships will require the ability to test hypotheses by running experiments on the physical system.

2. **Model construction with robust and efficient AI and ML to create the DT**. The core of the DT is a model that is built and updated based on data from the physical system and the environment in which it operates. These systems usually include multiple spatial and temporal scales and multimodal data involving large quantities, many of which may not be measured directly. Furthermore, the process of creating the model naturally includes significant challenges in data reduction, which focuses on identifying and representing the information contained in data. Challenges include identifying the model form, training data, training process, and the appropriate a priori information needed to construct the DT.

3. **Continuous learning**. The very nature of the DT requires that the physical system be monitored continuously and that this data be used to update the DT to reflect the current state of the system with the recognition that the state of the system drifts. Challenges include selecting the data appropriate for continuous learning, learning in an online mode in which data can be used only once at the time of collection, and being confident that the training algorithms *forget* when appropriate.

4. **Anomaly detection**. Anomaly detection can be described as the ability to identify or predict system behaviors or environments that were not part of the data used for designing or training the DT. This can include system failures as well as adversarial attacks against the physical system (as opposed to adversarial attacks against the training and operation of the DT). This is a critical part of any assurance effort but is identified as a separate area here because of the challenge of detecting and identifying system states that do not appear in training data.

5. **Codesigned software and hardware ecosystem**. There are two aspects to this challenge. First, the physical twin must be engineered to interact with the DT. This includes sensor design and placement, power management, and incorporating edge-based computational capabilities and control systems that can interface with the DT. This also includes designing the capability to deal with robustness and resilience issues, including data issues and adversarial attacks that are introduced through a DT. Second, the DT itself introduces hardware and software ecosystem challenges, including communication and bandwidth challenges, and the challenge of federation when many copies of a physical system—each with an individualized DT—are deployed.

6. **Standardization and metrics**. The development and deployment of DTs are currently very system and application specific. As DTs become common in engineered systems, standard protocols for their design, production, deployment, and maintenance will become necessary. In particular, safety issues that arise when using DTs will drive a regulatory environment that will require standardization. In addition to being necessary for safety and regulatory purposes, standardization and metrics will improve interoperability and performance and enable an overall increase in efficiency in the design and engineering of DTs.

Future AIRES workshops will continue to examine the foundation aspect of DTs. Each workshop is expected to address one or more of these specific research challenges in addition to looking at the broader issues associated with the development and deployment of DTs.

APPENDIX A. AIRES 2 WORKSHOP BREAKOUT REPORTS

Section	Title	Lead(s)
A.1	Machine Learning for Robust Digital Twins	Nathaniel Trask
A.2	Requirements and Methods for DT Construction	Samrat Chatterjee
A.3	Time Series Prediction for Digital Twins	Frank Liu
A.4	Robustness and Validation of Model and Digital Twins Deployment	David Stracuzzi Jenifer Shafer
A.5	Digital Twins for Real-Time Control Systems	Christine Sweeney
A.6	Methods for Continual and Online Learning for Digital Twins	Nurali Virani
A.7	Constructing Digital Twins using High-Dimensional Data	Arvind Mohan
A.8	Constructing and Using Digital Twins for Anomaly Detection	Adi Hanuka Jiaxin Zhang
A.9	Hardware and Software Issues in Edge Computing and Production Deployment of Digital Twins	Ron Oldfield
A.10	Scaling and Realization of Digital Twins on Cloud-and-HPC Systems	Piyush Modi
A.11	Nuclear Energy: Challenges and Applications of Digital Twins	Prashant Jain
A.12	Digital Twin Certified Additive Manufacturing	Aric Hagberg

A.1. MACHINE LEARNING FOR ROBUST DIGITAL TWINS

Chair:	Nathaniel Trask	Sandia National Laboratories (Sandia)
Participants:	John Emery	Sandia
	Hoang Tran	ORNL
	Patrick Blonigan	Sandia
	Ravi Raveendra	ESI Group
	Sanjay Choudhry	NVIDIA
	Guannan Zhang	ORNL
	David Schmidt	University of Massachusetts Amherst
	Felipe Viana	University of Central Florida
	Zhehui (Jeph) Wang	Los Alamos National Laboratory (LANL)

Introduction:

Digital twins (DTs) require a *digital thread* between a physical system and its simulated counterpart, which requires fast-forward simulations for the twin to provide predictions of the system and fast data assimilation tools to extract an appropriate model for the twin. Both models must be reasonably close to real-time to allow online inference.

Machine learning (ML) tools, and particularly deep learning (DL), have gained attention as potential means of supporting both fast partial differential equation (PDE) discretizations and learning models from data, owing to their ability to handle high-dimensional data in a relatively black-box manner, with many examples in the literature of orders-of-magnitude speedup vs. traditional finite element/finite volume techniques. Although these tools form an attractive candidate for the digital thread, robustness guarantees and trusted AI are needed to guarantee numerical stability, physical realizability, out-of-distribution inference, and accuracy—particularly in the small data limits often encountered in science and engineering applications.

Physics-informed ML (physML) has emerged as a discipline in which the traditional tools from scientific computing and numerical analysis may be applied to impart prior physical and mathematical knowledge on the learning process. Many techniques seek to regularize a training loss with a PDE residual to penalize deviations from prior knowledge. These techniques have provided exciting first examples in which ML provides transformative predictions beyond the means currently available in traditional numerical techniques.

Although promising, these techniques lack the mathematical foundations of traditional forward simulation, precluding in some cases their application in high-risk/high-consequence engineering environments that are particularly relevant to the US Department of Energy (DOE). To serve as a reliable digital thread, additional research must establish rigorous guarantees so that data-driven models and fast surrogates maintain requisite trust and usability.

Guiding questions:

- The extrapolation problem: how do networks perform inference out of training distribution?
- Incomplete physics: sometimes the model form is only partially known and lacks complete equations. How does one parameterize unknown physics and move beyond parameter estimation?

- Training/optimization tools specifically for physML: what can be leveraged in science and engineering scenarios that have more exploitable structure than classical ML training?

- Data: small data constraints are the norm for expensive physical systems, which leads to overfitting and challenges with generalization. What synthetic/experimental data sets are available in the labs that could benefit the community? The lack of open-source data sets and benchmarks makes it difficult to reliably establish best practices and the state of the art.

- Legacy codes and nonintrusive ML: how does one integrate the historical simulation codes that are ubiquitous in DOE but not equipped with the automatic differentiation required for many ML approaches?

- How does one move beyond TensorFlow/PyTorch and scale up to the large problems and complex geometries representative of DOE problems? How does one incorporate MPI (Message Passing Interface) parallelism when the data science industry is focused on the graphics processing unit (GPU) and the tensor processing unit?

- Lack of common frameworks for developing software: how does one mitigate duplicated effort and reinventing the wheel?

Key challenges:

- How does one guarantee convergence and stability of physics-informed surrogates and move beyond *eyeball norm* predictions?

- Unlike traditional (e.g., finite element) simulation, deep-learned surrogates lack convergence guarantees, which renders mesh refinement studies impossible. Recent approximation theory proves that networks are capable of better approximation than finite element/volume/difference codes, but can these be provided reliably in practice?

- Current physics-informed techniques only impose physics by penalty, which leads to predictions with on the order of 1%–0.01% mismatch. Although acceptable for some applications, others need this to hold to machine precision. There are other issues related to numerical treatment of physics, such as how to impose boundary conditions and initial conditions.

- How can one compare methods and determine the best approach? The standardized benchmarks that are ubiquitous in the classical ML world are not available for physML. In other words, what could the MNIST (Modified National Institute of Standards and Technology database) of the physML world do to allow for clear comparisons of the new methods emerging every day?

- How are missing physics handled? Many first results in physML have estimated physical parameters, but how does one move to problems ubiquitous in multiscale/multiphysics for which closures of unknown form must be found from data? Can physical realizability be obtained from models without applying dictionary learning?

State of the art:

- Currently available physML tools provide robustness in small data limits, which is critical for building DTs of engineered systems and improving extrapolation/out-of-distribution inference. Current tools are particularly promising/useful for inverse/optimization problems and physics discovery beyond forward simulation.

- PhysML is a young field, but several methods have emerged: physics-informed neural networks (PINNs), data-driven reduced-order models, Kutz's SINDy (sparse identification of nonlinear dynamics) method and related dictionary-based techniques, and operator regression.

- The roadblocks for these techniques are a lack of trustworthiness regarding guaranteed performance—generally, methods require hyperparameter optimization that still needs a human in the loop, which is undesirable for DTs. There is also a lack of formal error analysis and means to establish accuracy comparable to traditional forward simulation.

Research opportunities:

- Research in the past few years has focused on establishing initial proof-of-concept and feasibility of physML. In the coming years, there is a major opportunity for establishing broadly adopted benchmarks for pursuing more challenging metrics beyond the eyeball norm for physical systems. For example, establishing a standard problem could demonstrate how approaches converge with respect to data and architecture size and provide higher-order metrics of physically relevant statistics (e. g., demonstrating for turbulence—not just mean velocity and energy spectra but two-point statistics).

- Common platforms are necessary for establishing physics-informed learning best practices, so that physML can move beyond academic/hero problems and heuristics to obtain well-understood workflows applicable in industrial settings.

- Leverage broad opportunities to apply the design principles and expertise related to traditional numerical analysis and approximation theory. For example, most PINN approaches apply point collocation PDE residuals, while more recent work applies more advanced PDE discretization concepts. Mathematicians have been establishing the approximation theory properties of deep networks—can these be realized in practical settings? Many of these skill sets are ubiquitous throughout DOE, and many opportunities exist for repurposing non-data scientists to apply advanced optimization, solvers, and discretization expertise to improve physML.

Impact if research opportunities are addressed:

- If physML can be developed to be as trustworthy as PDE/ODE (ordinary differential equation) discretization tools, it is an ideal candidate for providing the beyond forward simulation tasks and data assimilation necessary for DTs.

- Existing technology does not close the loop efficiently. Potential improved performance promised by physML could allow for onboard smart devices with real-time sensing and feedback.

- Engineering and physics systems have a higher threshold for reliability and trustworthiness than the ML counterparts. Establishing more accurate physics-informed tools will translate to DTs that can identify potential gains in performance and more reliable workflows.

Summary:

DTs require efficient and reliable numerical tools to close the loop and provide fast prediction and assimilation of a given system. Current modeling and simulation tools are too costly to close this gap. PhysML provides a means to exploit the efficiency of ML while providing some of the robustness and reliability of physics-based models.

Although promising, these models have weaknesses compared to the finite element method and the finite volume method models serving as the workhorses for science and engineering problems. The reliability of physML methods must be established before they can play a critical role in the unsupervised tasks necessary for DTs.

A.2. REQUIREMENTS AND METHODS FOR DT CONSTRUCTION

Chair:	Samrat Chatterjee	Pacific Northwest National Laboratory (PNNL)
Cochair:	Massimiliano (Max) Lupo-Pasini	ORNL
Participants:	Thomas Britton	Jefferson Laboratory
	Bill Spotz	DOE
	Kristopher Velazquez	Lockheed Martin
	Ravi Raveendra	ESI Group
	Jibonananda Sanyal	ORNL

Introduction:

DTs are gaining increasing interest in the scientific community as virtual tools that provide an abstract representation of complex environments. The use of DTs in science would benefit experimental design and anomaly and failure detection/prediction to attain a certified self-reliance via an autonomous and continuous safety assessment.

PhysML was recently integrated with existing domain specific applications. However, the construction paradigms, safety requirements, and protocols that follow DT production and deployment depend on the application of interest. This hinders an objective verification and validation of the DT performance and objective comparison of different DTs for the same goal.

Thus, standardized terminology, protocols, and requirements are needed in the DT community to elevate the construction and performance evaluation of a DT above the details of a specific application.

Guiding questions:

- Are there standard guidelines to construct a DT that can be applied across different disciplines?
- Are there safety checks recognized by the IoT community to guarantee that DTs are self-reliant after deployment?
- Are there common paradigms that every DT is constrained to respect?
- Are there benchmarks that are recognized in the DT community as valid pass/fail tests that certify the correct functioning of the DT?

Key challenges:

- Standardization, generalizability, and universality of DT requirements. DTs are application driven, and thus their construction and performance assessment are bound to the specifics of the given application domain. This results in a heterogeneous nomenclature, list of necessary and sufficient conditions, and different metrics for success across the applications, which hinders scientific advance in the DT community as a cross-cutting discipline. A standard terminology would benefit future research opportunities and collaborations.
- Capture and address the complexity of environments. Developers should ask whether they need a unique DT or an ensemble of DTs with each DT focused on more specific tasks.

- Methods to certify DTs, including self-sustainability. Just as humans can reach success through several failures, the role of good-faith failure in the evaluation of the DT's self-reliance must be determined.

State of the art:

- Specific applications have made promising progress using hybrid and ML methods.

- ML and physics are currently integrated to generate an established groundwork by replacing agnostic ML models with more specific domain-driven physics models. However, DT construction and deployment is still very task specific.

- Application-specific standards are emerging but are not readily transferable to blueprints, processes, and best practices across applications.

- Modules are robust when used alone, but they still struggle as a cohesive system of modules. Integrating different DTs to generate an autonomous system of DTs raises challenges that are not entirely addressed (i.e., every component functioning properly as a standalone DT is not sufficient to guarantee that the whole system will work once the DTs are combined).

- The real world is complex, dynamic, and uncertain, which creates challenges. There are multiple layers of complexity: a single task can already be complex in itself; on top of that, this complex task may need to be combined with more complex realities (e.g., multiagent environment) in which multiple DTs interact with each other and environmental variables that were not accounted for in the task-learning phase.

Research opportunities:

- Develop modular approaches (System of Systems) for DT construction. This includes modularity in the design stage, instead of just brute force modularity during execution. This can include object-oriented paradigms used in programming (e.g., classes of classes, classes with multiple methods, one class construct can be used to characterize different things) to allow customization but still preserve a uniform/cross-cut paradigm. This would also favor integration of DTs that share the same structure, and it would reduce design costs.

- Automated awareness with uncertainty. DTs must distinguish between anomalous behavior and the natural evolution of a system (e.g., a new emerging paradigm). They must discriminate and understand how to handle these different scenarios autonomously within allowed error.

- Strengthen connections to meta learning (i.e., learning to learn paradigm), continual learning (i.e., incorporate new scenarios in the range of data previously explored), and federated learning (i.e., multiple replicase of the same DT are exposed to different scenarios, and the info is virtually shared via an outsourced server).

- Create clear standardized definition of rules of engagement. Engineers must understand advantages and disadvantages between multiagent DTs (i.e., multiple versions of DTs doing different things for different goals) and hierarchical DTs (i.e., the top-level agent monitors the common goal, and the goal/task is broken down into subtasks, and each DT focuses on subtasks).

- Create clear standardized definitions of success to determine that the DT is self-reliant and ready for deployment. Although different applications lead to different standards, a uniform protocol will

benefit transparency in terms of reproducibility of results. Standard and quantitative metrics that help measure success of a DT independent of the application field are also needed.

- For example, mean-square error or F-score in statistics are very broad and still used to assess the predictive performance of ML models independent of the application domain. Similar measures are needed for DT.

- Benchmark data sets that can be used to compare the performance of different DTs and measure progress and improvements of DTs are also needed.

- Certified protocols should be created to ensure the reliability of a DT throughout its lifespan. Reliable DT performance one instance is not sufficient for credibility, especially when modelling complex phenomena.

- Documentation standards—including a full history of failures, a clear explanation of how those failures were addressed during the DT design and deployment, and when self-learning and self-correction was needed—should be investigated.

Impact if research opportunities are addressed:

- Heterogenous operations, including the learning of new and more complex skills with computational agents and transferability to learn a set of skills.

- Outperforming current task learning with skill-based learning to perform multiple complex tasks.

- Certifiable procedures and processes to assess performance, including reliability and robustness of DT, including repositories to record compliance.

Summary:

A cross-cut standardization is needed for the creation, deployment, and self-sustainability process that evaluates the performance of a DT throughout its lifespan. If attained, this would foster scientific advancement through a more efficient integration of DT models with existing scientific capabilities.

Moreover, the community must develop a standard paradigm that lists the essential features that a DT must have, and this standard paradigm should be recognized and followed by industries and research institutes. Standardized terminology, procedures, and certification protocols would guarantee a clear reproducibility of existing results, and this would in turn foster future research and scientific progress.

Supporting Material:

Papers:

[1] David Jones, Chris Snider, Aydin Nassehi, Jason Yon, Ben Hicks, "Characterizing the Digital Twin: A systematic literature review," *CIRP Journal of Manufacturing Science and Technology*, Volume 29, Part A (2020): pp. 36–52, https://doi.org/10.1016/j.cirpj.2020.02.002.

[2] J. Moyne et al., "A Requirements Driven Digital Twin Framework: Specification and Opportunities," in *IEEE Access*, vol. 8 (2020): pp. 107,781–107,801, https://doi.org/10.1109/ACCESS.2020.3000437.

[3] F. Laamarti, H. F. Badawi, Y. Ding, F. Arafsha, B. Hafidh and A. E. Saddik, "An ISO/IEEE 11073 Standardized Digital Twin Framework for Health and Well-Being in Smart Cities," in *IEEE Access*, vol. 8 (2020): pp. 105950–105961, https://doi.org/10.1109/ACCESS.2020.2999871.

[4] Guo, J., Zhao, N., Sun, L. et al, "Modular based flexible digital twin for factory design," in *J Ambient Intell Human Comput* 10 (2019): pp. 1189–1200, https://doi.org/10.1007/s12652-018-0953-6.

[5] M. Thomas, B. Klenz, and P. R. Goodwin, "Artificial and Human Intelligence with Digital Twins," IIC Journal of Innovation (November 2019), https://www.iiconsortium.org/news/joi-articles/2019-November-JoI-Artificial-and-Human-Intelligence-with-Digital-Twins.pdf.

URLs:

[1] Internet of Business, "Gartner: Four Best Practices for Managing Digital Twins," https://internetofbusiness.com/half-of-businesses-with-iot-projects-planning-to-use-digital-twin/

[2] Rachel Roundy, "Best Practices for Digital Twin Implementation," *DZone*, August 5, 2020, https://dzone.com/articles/best-practices-for-digital-twin-implementation-1.

[3] Michael Thomas, Brad Klenz, and Prairie Rose Goodwin, "How to use human and artificial intelligence with digital twins," *Control Engineering*, October 13, 2020, https://www.controleng.com/articles/how-to-use-human-and-artificial-intelligence-with-digital-twins/.

A.3. TIME SERIES PREDICTION FOR DIGITAL TWINS

Chair:	Frank Liu	ORNL
Cochairs:	Jan Drgona	PNNL
	Chetan Kulkarni	NASA Ames Research Center
Participants:	Willem Blokland	ORNL
	Alessandro Cattaneo	LANL
	Reese Jones	Sandia
	David Mascarenas	LANL
	Abhinav Saxena	GE Research

Introduction:

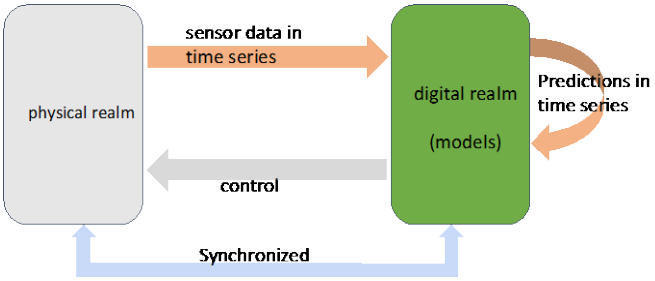

Time series analysis plays a critical role in the development and deployment of DTs. As a more general notion of the DT, the models in the digital realm should be synchronized with the corresponding system in the physical realm. The dynamic behavior of a physical system is often described by streams of time series data collected at different rates with different fidelity. Time series analysis includes parsing and preprocessing the time series data, constructing and fine-tuning system dynamics models in the digital realm, and facilitating the design and tuning of the controller when active controls are needed. All these tasks must be carried out with the consideration of sensor noises and uncertainties, which includes model parameter mismatch between the digital model and the corresponding physical system as well as unmodelled physics in the digital model.

The participants of this breakout session came from various DOE laboratories, NASA centers, and industry with broad backgrounds in applied math, ML/AI, industrial control, and high-energy physics. The lively discussions covered diverse topics on the state of the art in time series predictions, challenges, and the future research opportunities.

Key challenges:

- Quality and fidelity of time series data as the sensor input. Challenges include the explosive data size, the nonstationary nature of the data, and the sample rates required to capture the dynamic behavior of the physics realm, different modalities, and the jitter—especially for high-frequency sampling commonly used in accelerator electronics.

- The lack of a mature theoretical framework, especially for diverse problems encountered in DTs.

- In the same vein, there is the lack of a general application methodology (i.e., a unified method to deploy various ML techniques on problems). Most of the current approaches are based on trial-and-error and ad-hoc methods.

- A closely related challenge is how to incorporate domain knowledge in the time series prediction for DTs, such as causality discovery, interpretability, constraints, and symmetries.

- In some applications (e.g., accelerator physics, battery health prediction) the time series data demonstrates multiple time scales with orders of magnitude differences. Big challenges lie in how to store the large amounts of data on the theoretical framework for analysis and prediction.

- In a broader sense, big challenges lie ahead in persuading the broader ML community, which is mostly focused on natural language processing (NLP) problems and applications in financial markets, of the importance of time series prediction for DTs.

- The weak link from the time series predictions in the digital realm to the performance in the physical realm and back-to-the-model update. This feedback loop may also suffer long delays and was designed in a manual, ad-hoc fashion.

- Better understanding of a family of generative models with theoretical guarantees on stability and robustness against noise and adversarial attacks.

- Dealing with time-varying and delayed dynamical systems and distribution shifts.

State of the art:

- The participants concluded that the theoretical foundations are fragmented. There are several active research areas that can contribute to the time series prediction of DTs. The examples include differential equations of known physics, deep recurrent neural networks (RNNs), long short-term memory (LSTM), Neural ODE, and deepOpnets. Other examples include various autoregression models, state-space models, and probabilistic graph models. Some research advancements include domain decomposition methods and Koopman operators.

- The participants also concluded that the programming environments for time series prediction do not have a common reference platform. On one hand, there are popular ML and statistical analysis environments such as Tensorflow, PyTorch, and R. On the other hand, there are domain-specific software environments such as Matlab/Simlink, Simpscape, Modelica, Ansys's AGI software, and COMSOL Multiphysics models.

- Methods for dealing with multiscale physics are mostly ad-hoc.

Research opportunities:

- Leverage rich body of existing theories (e.g., control theory, dynamic systems, information theory) to advance a more rigorous and coherent theoretical framework.

- Hybrid models combining physics and data-driven methods.

- Learning-based control to use time series prediction to facilitate and improve design and tuning of controllers.

- For data processing, DT-aware data processing, especially the knowledge from domain experts and subject matter experts.

- Close integration of time series prediction tools and physical simulation environments. Additional research to incorporate control can be also fruitful.

- Coherent method to deal with time series prediction for multi-timescale physics.

Impact if research opportunities are addressed:

- More effective systems, health management, and prognostics framework with broader impact on maintenance and decision support.

- Computationally efficient DTs with performance guarantees, the capability for online continual learning, quantified uncertainty, and increased level of autonomy.

- Another benefit is better industrial-grade controllers, which are easy to use and easy to tune (e.g., classic PID [proportional–integral–derivative controller] controllers).

Summary:

DTs pose unique challenges and opportunities for time series research. Some unique research questions, such as physics-informed modeling and effective analysis of multiscale time series data, are not actively pursued by the broader ML community. It is highly desirable to engage the community with a clear problem definition and publicly available data sets.

On a longer time horizon, time series analysis and prediction requires successful development, maintenance, and deployment of analysis software—both at the centralized computing facilities and at the edge. Given the considerable cost of developing and maintaining a software ecosystem, it is also beneficial to explore and provide seed funding to mitigate the currently fragmented software ecosystem.

Supporting Material:

Papers:

[1] J. Guo, N. Zhao, L. Sun et al., "Modular based flexible digital twin for factory design" *J Ambient Intell Human Comput* 10 (2019): pp. 1189–1200, https://doi.org/10.1007/s12652-018-0953-6.

[2] Lu Lu, Pengzhan Jin, and George Em Karniadakis, "Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators," arXiv preprintarXiv:1910.03193, 2019.

[3] Maziar Raissi. "Deep hidden physics models: Deep learning of nonlinear partial differential equations," in *Journal of Machine Learning Research* 19, no. 1 (2018): pp. 932–955.

[4] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar, "Fourier neural operator for parametric partial differential equations," arXivpreprint arXiv:2010.08895, 2020.

[5] Gitta Kutyniok, Philipp Petersen, Mones Raslan, and Reinhold Schneider. "A theoretical analysis of deepneural networks and parametric PDEs," arXiv preprint arXiv:1904.00377, 2019.

[6] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud, "Neural ordinary differential equations," in *Advances in Neural Information Processing Systems* (2018):pp. 6571–6583.

[7] Jonas Adler and Ozan Öktem, "Solving ill-posed inverse problems using iterative deep neural networks," in *Inverse Problems*, 33, no. 12 (2019): 124007.

[8] Jan Drgona, Elliott Skomski, Soumya Vasisht, Aaron Tuor, and Draguna Vrabie, “Spectral Analysis and Stability of Deep Neural Dynamics,” arXiv:2011.13492, 2020.

A.4. ROBUSTNESS AND VALIDATION OF MODEL AND DIGITAL TWINS DEPLOYMENT

Chairs:	David Stracuzzi	Sandia
	Jenifer Shafer	DOE/Advanced Research Projects Agency–Energy (ARPA-E)
Cochairs:	Svitlana Volkova	PNNL
	Jaideep Ray	Sandia
Participants:	Matt Barone	Sandia
	Sharlotte Kramer	Sandia
	Daniel Ratner	SLAC National Accelerator Laboratory
	Maria Glenski	PNNL
	Andy Huang	Sandia
	Laura Pullum	ORNL

Introduction:

For DTs to become a central fixture in mission-critical systems, a better understanding is required of potential modes of failure, quantification of uncertainty, and the ability to explain a model's behavior. These aspects are particularly important because the performance of a DT will evolve during model development and deployment for real-world operations. The requirements for safety-critical DT systems—robustness, accountability, transparency, and fairness—are defined below:

- **Robust** DTs must be resilient to variations in data inputs.
- **Accountable** DTs must demonstrate reliability when applied in key circumstances and be able to review historical predictions and inferences (closely related to transparency).
- **Fair** DTs must be equitable across representative subsets (e.g., across subpopulations of users impacted by DT outputs).
- **Transparent** DTs must enable high-quality and correct interpretations of model behavior to identify points of failure through data inputs and model predictions. This can be accomplished either by interactive explanations of model behavior or by quantifying the predictive performance of the model. An example of DTs could include a combination of uncertainty quantification (UQ), out-of-distribution analysis, and traditional performance analysis (e.g., cross-validation).

A rigorous approach to DT validation and verification (V&V) is also required for three reasons: (1) to provide a basis for trust in DT adoption, (2) to reduce risk of DT backsliding, and (3) to decrease likelihood of accidents with a DT. DT V&V is challenging because testing standards depend on the area of application and the specifics of the ML approach used in the development of the DT.

Guiding questions:

- How is robustness defined?
- How is validation defined?

- How are current definitions of robustness and validation insufficient, and how could they be improved?

- What does it mean to do V&V on a DT model when some parts of the model are represented by ML (statistical) models whereas other parts are represented by mathematical (theoretical) models?

- Given that ML models extrapolate statistically rather than via an explicit encoding of underlying physics, what are the implications for DT model development, verification, and validation?

- What terms and methods, which are typically well-defined with respect to V&V for (equation-based) physics models, must be changed for ML models?

- How does the problem change for DTs that are updated over time (i.e., DT models that are incrementally revised as data is collected from real systems)?

- What are the near-term research needs to develop appropriate V&V methods?

Key challenges:

- Many fundamental questions about DT V&V and robustness remain open. For example, broadly accepted definitions of verification, validation, and robustness for DTs as well as for ML do not exist. The definition of V&V (i.e., safety, fairness, interpretability) varies between communities and provides constraints on solution space.

- V&V is a multidimensional process that must account for robustness (including extrapolation), fairness, interpretability, safety, and so on.

- Currently, it is not clear what metrics are important for DT V&V. Moreover, quantitative metrics might be insufficient; they must be paired with qualitative evaluation. What is the testing cadence and stopping criteria for DT V&V?

State of the art:

- Focus of DOE effort is on DT for high-consequence (aka safety-critical) systems – without effective V&V the adoption will not happen.

- The literature is extensive for classical V&V and UQ, and the literature is growing for ML evaluation, including recent work in adversarial attacks, input perturbation, UQ for some ML models, out-of-distribution, and training optimization, as described in detail below.

Related work that extends the traditional evaluation of ML models to address the issues of robustness, accountability, fairness, or transparency of model performance consider these issues independently, as summarized in Table 1. DT evaluation will require the integration of multiple packages or standalone tools with distinct requirements, environments, or interfaces to perform a multifaceted evaluation with tools that often face similar limitations. For example, existing tools that address model robustness are largely focused on adversarial attacks (e.g., Adversarial Robustness Toolbox [ART], Advertorch, Foolbox, Advbox, OpenAttack, TextFooler) or gradient attacks (e.g., Foolbox). Others are also tied to a specific architecture framework, such as PyTorch (e.g., Advertorch), which limits the flexibility when applied to previously or independently developed models.

Table 1. Summary of existing tools for ML evaluation highlighting related work across each of the four dimensions of interest

Robustness	Accountability	Fairness	Transparency
ART [20]	**GLUE** [31]	**Aequitas** [26]	**interpretML** [21]
Advertorch [7]	**SuperGLUE** [30]	**AI Fairness 360** [4]	**SHAP** [13]
Foolbox [24]	**SquAD 2.0** [23]	**Fairlearn** [5]	**Captum** [12]
Advbox [8]	**ROAR** [10]	**Fairness Indicators** [34]	**LIT** [29]
OpenAttack [33]		**FairSight** [1]	**TreeExplainer** [14]
TEAPOT [16]		**ML-fairness-gym** [35]	
TextAttack [19]		**Scikit-Fairness** [36]	
TextFooler [11]			

Roadblocks:

Coordinated focus of effort and funding for fundamental research.

Research opportunities:

- Develop broadly accepted, operational definitions of DT robustness, verification, validation, and fairness. The community must converge on fundamental definitions for validation, verification, robustness, and fairness in the context of DT in addition to specific applications and users. For example, in robust DT, outputs do not change drastically if the inputs are changed a little. DTs fail gracefully if fed nonphysical inputs.

- Validation has no change in goal because the DT still needs to reproduce physics, though the associated techniques may change. Verification may entail decoupling the DT into the PDE component and the ML component. The PDE component would be tested just as one does today and then compared with exact solutions and discretization error convergence. The ML component could be attacked with novel methods that go beyond Local Interpretable Model-Agnostic Explanations (LIME) and other 1D evaluation approaches.

- Develop qualitative approaches and quantitative metrics for evaluating DT robustness, fairness, interpretability, and so on. For example, robustness can be tackled with sensitivity analysis and validation with forward UQ, nominally, but Bayesian inverse problems may be solved to infer input uncertainties from other experiments. For verification, the ML piece must have at least been cross-validated, which is standard practice, but it does not speak to predictive uncertainty, extrapolation ability, or safety in the field. If the ML piece is a representation of training data from high-fidelity simulations, explanations must be generated for the ML model (e.g., via Deep Taylor Series or Generalized Additive Models [collectively called Locally Interpretable Model-agnostic Explanations]), and it must behave something like the high-fidelity simulations. This appears to be a brand new topic with no publications in this area.

- Fundamental research in UQ mathematics is required to support DT V&V. Major research questions include but are not limited to the following:

    - What requirements are placed on UQ for ML?
    - What constitutes a sufficient UQ evaluation of an ML model or an ensemble? There are many sources of uncertainty in an ML model; a complete evaluation is at least as intractable as it is for equation-based models.

- How does one combine/propagate UQ for ML components with physics equations? Forward UQ, nominally, but Bayesian inverse problems may be solved to infer input uncertainties from other experiments.

- Fundamental research in formal methods for evaluating critical model parts and scaling these formal methods is needed.

- Redesign of real-world systems to facilitate ML V&V and model development (e.g., data instrumentation).

Impact if research opportunities are addressed:

- What new scientific capabilities will be enabled?
  - Data-driven insights into physics/engineering theory
  - Precise analysis/quantification of model applicability
  - Model-driven automation of critical applications

- What will the new methods and techniques enable?
  - Certification, validation, and robustness methods and evaluation techniques
  - Data-driven models that admit a physics-based chain of reasoning for SciML models
  - Improved accuracy of scientific models without a significant increase in computational cost

Summary:

Targeted continuous evaluation of DT is required to estimate the tradeoff between the risks and benefits of deploying DT models. To achieve this goal, the community must admit that the current ML evaluation process is broken [37], and because the ML component is a critical part of DT, fundamentally novel approaches for evaluating mission-critical DT systems are needed.

The ubiquitous method of multifold train and validate followed by the evaluation on a final holdout test set assumes that the data is sampled from a distribution that represents the data that the model will see after deployment and during the operational phase. For several reasons, this assumption often does not hold. The objective of the model developer is often to evaluate the model performance, optimize hyperparameters, and instill the developer with confidence in the model's performance. The focus is not on identifying as many modes of model failure as possible and determining how best to correct or mitigate them, but it should be.

Supporting Material

Papers:

[1] Y. Ahn and Y. R. Lin, "Fairsight: Visual analytics for fairness in decision making," in *IEEE transactions on visualization and computer graphics* (2019).

[2] D. L. Arendt, "Parallel embeddings: a visualization technique for contrasting learned representations," in *Proceedings of the 25th International Conference on Intelligent User Interfaces* (2020).

[3] D. Arendt, Z. Huang, P. Shrestha, E. Ayton, M. Glenski, and S. Volkova, "CrossCheck: Rapid, Reproducible, and Interpretable Model Evaluation," arXiv preprint (2020).

[4] R. K. Bellamy, K. Dey, M. Hind, S. C. Hoffman, S. Houde, K. Kannan, et al., "AI Fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias," in *IBM Journal of Research and Development* (2019).

[5] S. Bird, M. Dudik, R. Edgar, B. Horn, R. Lutz, V. Milan, M. Sameki, H. Wallach, and K. Walker, *Fairlearn: A toolkit for assessing and improving fairness in AI*, Technical Report MSR-TR-2020-32 (Microsoft, May 2020).

[6] A. R. Chouldechova, "A snapshot of the frontiers of fairness in machine learning," in *Communications of the ACM* 63, no. 5 (2020).

[7] G. W. Ding, L. Wang, and X. Jin, "AdverTorch v0.1: An Adversarial Robustness Toolbox based on PyTorch," arXiv:1902.07623 (2019).

[8] D. Goodman, H. Xin, W. Yang, W. Yuesheng, X. Junfeng, and Z. Huan, "Advbox: a toolbox to generate adversarial examples that fool neural networks," arXiv preprint (2020).

[9] B. A. Haibe-Kains, "Transparency and reproducibility in artificial intelligence," *Nature* 586, E14–E16 (2020).

[10] S. Hooker, D. Erhan, P. J. Kindermans, and B. Kim, "A benchmark for interpretability methods in deep neural networks," in *Advances in Neural Information Processing Systems* (2019).

[11] D. Jin, Z. Jin, J. T. Zhou, and P. Szolovits, "Is BERT Really Robust? Natural Language Attack on Text Classification and Entailment," arXiv preprint (2019).

[12] N. Kokhlikyan, V. Miglani, M. Martin, E. Wang, B. Alsallakh, J. Reynolds et al., "Captum: A unified and generic model interpretability library for PyTorch," arXiv preprint (2020).

[13] S. M. Lundberg and S. I. Lee, "A Unified Approach to Interpreting Model Predictions," in *Advances in Neural Information Processing Systems* 30 (2017).

[14] S. M. Lundberg, G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, B. Nair, and S. I. Lee, "Explainable AI for Trees: From Local Explanations to Global Understanding," arXiv preprint (2019).

[15] N. M. Mehrabi, "A survey on bias and fairness in machine learning," arXiv:1908.09635 (2019).

[16] P. Michel, G. Neubig, X. Li, and J. M. Pino, "On Evaluation of Adversarial Perturbations for Sequence-to-Sequence Models," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics* (2019).

[17] A. Modas, S. Moosavi-Dezfooli, and P. Frossard, "Sparsefool: A Few Pixels Make a Big Difference," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2019).

[18] C. Molnar, *Interpretable Machine Learning* (2019).

[19] J. X. Morris, E. Lifland, J. Y. Yoo, J. Grigsby, D. Jin, and Y. Qi, "TextAttack: A Framework for Adversarial Attacks, Data Augmentation, and Adversarial Training in NLP," arXiv preprint (2020).

[20] M.-I. Nicolae, M. Sinn, M. N. Tran, B. Buesser, A. Rawat, M. Wistuba et al., "Adversarial Robustness Toolbox v1.0.0," arXiv preprint (2018).

[21] H. Nori, S. Jenkins, P. Koch, and R. Caruana, "InterpretML: A Unified Framework for Machine Learning Interpretability," arXiv preprint (2019).

[22] D. Pruthi, "Combating Adversarial Misspellings with Robust Word Recognition" (The 57th Annual Meeting of the Association for Computational Linguistics [ACL], 2019).

[23] P. Rajpurkar, R. Jia, and P. Liang, "Know What You Don't Know: Unanswerable Questions for SquAD," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)* (2018).

[24] J. Rauber, R. Zimmermann, M. Bethge, and W. Brendel, "Foolbox Native: Fast adversarial attacks to benchmark the robustness of machine learning models in PyTorch, TensorFlow, and JAX," in *Journal of Open Source Software* (2020).

[25] E. Saldanha, B. Praggastis, T. Billow, and D. Arendt, "ReLVis: Visual Analytics for Situational Awareness During Reinforcement Learning Experimentation," in *EuroVis (Short Papers)* (2019).

[26] P. Saleiro, "Aequitas: A Bias and Fairness Audit Toolkit," arXiv preprint (2018).

[27] S. H. Silva, "Opportunities and challenges in deep learning adversarial robustness: A survey," arXiv:2007.00753 (2020).

[28] D. J. Stracuzzi, M. G. Chen, M. C. Darling, M. G. Peterson, and C. Vollmer, *Uncertainty quantification for machine learning*, SAND2017-6776 (Sandia National Laboratories, 2017).

[29] I. Tenney, J. Wexler, J. Bastings, T. Bolukbasi, A. Coenen, S. Gehrmann, and A. Yuan, "The Language Interpretability Tool: Extensible, Interactive Visualizations and Analysis for NLP Models," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing* (2020).

[30] A. Wang, Y. Pruksachatkun, N. Nangia, A. Singh, J. Michael, F. Hill, and S. Bowman, "SuperGLUE: A stickier benchmark for general-purpose language understanding systems," in *Advances in Neural Information Processing Systems* (2019).

[31] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. Bowman, S., "GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding," in *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP* (2018).

[32] F. Yang, "How do visual explanations foster end users' appropriate trust in machine learning?" in *Proceedings of the 25th International Conference on Intelligent User Interfaces* (2020).

[33] G. Zeng, F. Qi, Q. Zhou, T. Zhang, B. Hou, Y. Zang, and M. Sun, "OpenAttack: An Open-Source Textual Adversarial Attack Toolkit," arXiv preprint (2020).

[34] "Fairness Indicators," GitHub repository, updated March 7, 2022, https://github.com/tensorflow/fairness-indicators.

[35] "ML-fairness-gym," GitHub repository, updated February 4, 2022, https://github.com/google/ml-fairness-gym.

[36] "Scikit-Fairness," GitHub repository, updated February 7, 2021, https://github.com/koaning/scikit-fairness.

[37] Will D. Heaven, "The way we train AI is fundamentally flawed," *The Technology Review*, November 18, 2020, https://www.technologyreview.com/2020/11/18/1012234/training-machine-learning-broken-real-world-heath-nlp-computer-vision/.

A.5. DIGITAL TWINS FOR REAL-TIME CONTROL SYSTEMS

Chair:	Christine Sweeney	LANL
Cochairs:	Sandra Biedron	Element Aero and University of New Mexico
	Malachi Schram	PNNL
Participants:	Mariana Fazio	University of New Mexico
	Salvador Sosa Guitron	University of New Mexico
	Hong Wang	ORNL
	Christopher Mayes	SLAC National Accelerator Laboratory
	Dave Caulton	Element Aero
	Huafeng Yu	Boeing

Introduction:

DTs are being increasingly developed for fast real-time control of complex and dynamic systems. DTs for this purpose have unique characteristics. Because they are deployed for real-time use, they must run fast and have their physical constraints met, such as being runnable where needed (possibly and including at the edge) with enough resources (e.g., power, bandwidth, computing performance, and storage). They may be part of a system of distributed control [6] or coordinated via network sensors, each of which may also have its own DT. They can consume data in real time for decision support and require adequate and relevant sensor data and sampling rates as well as somewhat predictable inputs. What is the correct balance of the number and type of sensors, and when is there enough data to make a reasonable prediction? They must be resilient in the event of data loss, corruption, variance, and/or uncertainty and may need backup DTs in the event of catastrophic failure. They also may be part of critical systems that cannot afford to use trial and error for control. The control systems that they belong to must have stopping conditions, and the DTs must help with detecting those.

The DT can be implemented as a fast-running physical simulation or as an ML-based surrogate (e.g., using neural networks or Gaussian processes). These DTs are often specialized to provide only the essential information to make control decisions—they are not general purpose. If the DT is ML based, it may be possible to pretrain them and transfer what they have learned to production use with minimal or no retraining, or continuous learning may be necessary. The control workflow may have human-in-the-loop as part of a control mechanism. Control systems that use DTs can use various forms of optimization, reinforcement learning (RL) [12], and adaptive control.

Use cases for DTs for control span many areas: transportation (e.g., vehicles, aircraft, spacecraft); facility operations and instrumentation including robotics; computation (e.g., load balancing, cybersecurity, model prediction); user-level data acquisition, assimilation and experimental or observational equipment control; and process control (e.g., chemical, nuclear, electrical, manufacturing).

Guiding questions:

- What are the defining challenges for DTs for real-time control systems?
- Where does the community stand today?
- Where does the community want to go?

- What are the requirements for DTs? For example, in RL or for prediction in the design of experiments?

- What are useful computational workflows for training, deploying, and updating DT-based control solutions using high-performance computing (HPC) resources and edge computing resources?

Key challenges:

- Data challenges: DTs for control have an added responsibility when considering that incoming input data may be missing, noisy, corrupt, or lacking sufficient metadata for comprehension. The state of input may be partially observable, which may limit the certainty of DT outputs. Incoming data may be multimodal to reflect the complex environment in which it is operating or observing. Errors in the incoming data may be from different sources: theory, knowledge, parameters, or sensors. Historical data may also not be available. Depending on the sector, there may not be comparable data available. For example, aircraft propulsion systems have many engines available of the same model/type that can enhance the data streams. This is not the case for bespoke machines for scientific discovery, but data could be the future for the development of near-autonomous systems such as radiation therapy.

- Model development/design challenges: An effective strategy may be to create DTs for control that rely on more than one source of information to provide outputs that can guide control, including training data, physics-based information, and/or uncertainties. However, this may add complexity to the DT. The DT may need to do continuous learning, which could pose challenges to learning at the edge or where resources may not be readily available. Cyberattacks can be made on DTs within autonomously controlled systems, and preventing those must be a priority.

- Deployment Challenges: Chief among deployment concerns discussed by the group were challenges with safety or performance guarantees as well as thorough validation and testing of DTs in control systems. Other deployment challenges include responding to changes in the physical system (e.g., a part is added or replaced), which may require rebuilding or modifying the DT and then testing and validating it again. For example, the large scale of coordination required in physical systems that have many sensors and require multiple controllers and/or DTs can be challenging.

- Safety assurance challenge: Owing to the introduction of new paradigms in model-based and learning-based designs for control systems, a major challenge is the lack of appropriate verification, validation, and safety assurance technologies. Most conventional testing, verification, and simulation technologies do not fit well because of their limited capability, performance, and trustworthiness in the new designs—particularly the learning-based design. Another challenge is the lack of good evaluation methods and metrics that can be used as proof to build trust and confidence in certification and deployment.

State of the art:

- In general, the group felt that the state of the art is still in the early days for DTs and control. It is a long-term goal at this point and depends on the maturation of DTs, real-time data analytics (i.e., to determine state), edge computing systems, availability of sensors/data, and many other factors to become a mature technology.

- Forays into DT-based control are being made, for example, in aviation, accelerators [1–6], and self-driving cars. Several use scenarios were identified: (1) online DTs model the physical system or environment and provide the controller with observations, (2) DT of an entire physical system used to

generate data and try different control strategies, and (3) smaller control loops are used to advise larger simulations.

- A sizable number of groups are working on and producing standards, especially in aviation. Standards are emerging for how to best test and validate these control systems and make them safe (see standards references below).

Research opportunities:

Several research opportunities for DTs exist in this area, and as noted earlier progress in control systems also depends on advances in components that make up the larger control workflow [7].

- Research is needed to make fast, hybrid or semi-physical models that are based both on data (e.g., trained neural network or Gaussian process models) and physics or computational models based on the domain. These hybrid models are needed to provide speedy and accurate predictions based on real-world input. Some use cases are complex enough to require multiple models within subsystems, so research into how to compose these and tools to aid in model composition are needed. Components may be at different scales and must be connected in ways that bridge scales. If uncertainties are used to provide robust control, they must propagate through a system of composed models. Research in how to represent uncertainties within such systems is needed for stochastic dynamics control in particular [8].

- Research is also needed in data analytics and ML to process observations input to the DT to translate these to unique states and to output unique states as well, which is critical for an accurate control system. In addition to unique states, research on building DTs that can operate on partial information is also important because not all observations will be complete in a control situation. Research in fuzzy logic and/or neural-fuzzy [9] will help with DTs that must operate using vague or imprecise observations. It will be useful to develop abstractions and ontology around the types of observations and conditions under which the DT for control would operate.

- Several scientific standards are emerging for autonomous control using DTs and will serve as some of the first of their kind for this field. Thorough revisions and reviews will be needed as these are developed because many of these system deployments hold significant risk if the systems malfunction. It is important that the standards address all aspects of testing, evaluation, and V&V, and novel methods will be required to do so. Data sets for bringing the systems to the standards must be collected or generated and made widely available. Ensuring that these data sets are representative of the domain will be a challenging area of research and development.

Impact if research opportunities are addressed:

Automatic control is key to reducing operational costs, improving efficiency, and improving data quality in many experimental facilities and systems. For example, expensive facilities such as particle accelerators can save minutes to hours on calibration and setup times. This would allow for additional experimental beam time as well as better data, which would better enable and accelerate scientific discovery. For transportation systems, automation can greatly improve safety and assurance because most errors are caused by humans. In the area of cybersecurity, automation can accomplish consistent vigilance over systems that are continually under attack, thereby ensuring security for valuable information.

Summary:

DTs used in real-time control of systems have unique requirements and challenges because they are part of critical systems that have related constraints and risks. Several implementation methods can be used for these DTs, including neural networks, emulators, and/or hybrid physics models. Challenges for DTs are that they run in real time; cope with partial, incomplete, vague, or corrupt data; may require input from many sources; and may have complex or challenging deployment constraints. Currently, a few systems are being developed for transportation and particle accelerators. Standards are being written to govern testing, verification, and validation of these systems. Research is needed in the development of hybrid DT models, data analytics to determine accurate state information, and standards development to ensure systems are safe. The impact of doing this research will result in faster scientific discovery, reduced operational costs, improved efficiency, and added safety and security.

Supporting Material:

Papers:

[1] C. Emma, A. Edelen, M. J. Hogan, B. O'Shea, G. White, and V. Yakimenko, "Machine learning-based longitudinal phase space prediction of particle accelerators" in *Phys. Rev. Accel. Beams* 21, 112802 (2018).

[2] C. Emma, A. Edelen, A. Hanuka, B. O'Shea, and A. Scheinker, "Virtual Diagnostic Suite for Electron Beam Prediction and Control at FACET-II," *Information* 12, no. 61, (2021): https://doi.org/10.20944/preprints202101.0115.v1.

[3] S. Hirlander and N. Bruchon, "Model-free and Bayesian Ensembling Model-based Deep Reinforcement Learning for Particle Accelerator Control Demonstrated on the FERMI FEL" (2020): https://doi.org/10.48550/arXiv.2012.09737.

[4] Verena Kain, Simon Hirlander, Brennan Goddard, Francesco Maria Velotti, Giovanni Zevi Della Porta, Niky Bruchon, and Gianluca Valentino, "Sample-efficient reinforcement learning for CERN accelerator control" in *Phys. Rev. Accel. Beams* 23, 124801 (December 2020).

[5] R. Roussel, A. Hanuka, and A. Edelen, "Multi-Objective Bayesian Optimization for Accelerator Tuning," arXiv:2010.09824 (2020): https://doi.org/10.48550/arXiv.2010.09824.

[6] J. St. John, G. Perdue, J. Duarte, M. Schram et al., "Real-time Artificial Intelligence for Accelerator Control: A Study at the Fermilab Booster," arXiv:2011.07371 (2021).

[7] L. Petersson, D. Austin, and H. Christenseni, "DCA: a distributed control architecture for robotics," *Proceedings of the 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Maui, HI, USA, 4 (2001): pp. 2361–2368, https://doi.org/10.1109/IROS.2001.976423.

[8] T. Chai, J. Qin, and H. Wang, "Optimal operational control for complex industrial processes," *Annual Reviews in Control*, 38 (2014): pp. 81–92.

[9] H. Wang, P. Afshar, ILC-based fixed-structure controller design for output PDF shaping in stochastic systems using LMI techniques, *IEEE Transactions on Automatic Control*, Vol. 54, No. 4, pp. 760–773 , 2009.

[10] J. Noriega and H. Wang, "A direct adaptive neural network control for unknown nonlinear systems and its application," *IEEE Transactions on Neural Networks* 9, (1998): pp. 27–34.

[11] Huafeng Yu, Xin Li, Richard Murray, Ramesh S, Claire Tomlin (Eds.), *Safe, Autonomous and Intelligent Vehicles* (Springer 2018), ISBN 978-3-319-97300-5, 2018.

[12] Steve Beland, Isaac Chang, Alex Chen, Matt Moser, Jim Paunicka, Doug Stuart, John Vian, Christina Westover, and Huafeng Yu, "Towards Assurance Evaluation of Autonomous Systems," (International Conference on Computer-Aided Design [ICCAD], 2020).

[13] R.S. Sutton and A.G. Barto, *Reinforcement learning: An introduction* (MIT press, 2018).

Reports:

[14] R. Stevens et al., *AI for Science*, US Department of Energy (2019): https://www.anl.gov/ai-for-science-report.

[15] *The National Artificial Intelligence Research and Development Strategic Plan*, National Science and Technology Council (2016): https://www.nitrd.gov/pubs/national_ai_rd_strategic_plan.pdf.

[16] *Roadmap for Intelligent Systems in Aerospace*, AIAA, Intelligent Systems Technical Committee (2016): https://studylib.net/doc/7217647/aiaa-roadmap-for-intelligent-systems-in-aerospace.

[17] *2015 NASA Technology Roadmaps*, *TA 4: Robotics and Autonomous Systems*, NASA (2015): https://www.nasa.gov/sites/default/files/atoms/files/2015_nasa_technology_roadmaps_ta_4_robotics_and_autonomous_systems_final.pdf

[18] *Autonomy Research for Civil Aviation: Toward a New Era of Flight*, National Research Council (2014): https://www.nap.edu/catalog/18815/autonomy-research-for-civil-aviation-toward-a-new-era-of.

[19] Executive Office of the President, "Fiscal Year (FY) 2022 Administration Research and Development Budget Priorities and Cross-cutting Actions," Memorandum for the Heads of Executive Departments and Agencies, US M-20-29, August 2022: p. 3, https://www.whitehouse.gov/wp-content/uploads/2020/08/M-20-29.pdf.

[20] US Department of Energy Office of Science, *Basic Research Needs Workshop on Compact Accelerators for Security and Medicine*, May 2019, https://science.osti.gov/-/media/hep/pdf/Reports/2020/CASM_WorkshopReport.pdf?la=en&hash=AEB0B318ED0436B1C5FF4EE0FDD6DEB84C2F15B2.

Standards:

[21] SAE International G-34/EUROCAE WG-114 Committee on Artificial Intelligence in Aviation, *Process Standard for Development and Certification/Approval of Aeronautical Safety-Related Products Implementing AI*, https://www.sae.org/works/documentHome.do?comtID=TEAG34&inputPage=wIpS.

[22] IEEE Standards Association, *Raising the Standards in Artificial Intelligence Systems (AIS)*, https://standards.ieee.org/initiatives/artificial-intelligence-systems/index.html.

[23] OCEANIS, *Global AI Standards Repository*, https://ethicsstandards.org/repository/.

[24] Axel Huebl, Remi Lehe, Jean-Luc Vay, David P. Grote, Ivo F. Sbalzarini, Stephan Kuschel, David Sagan, Christopher Mayes, Frederic Perez, Fabian Koller, and Michael Bussmann. "openPMD: A meta data standard for particle and mesh based data," (2015): https://doi.org/10.5281/zenodo.591699.

URLs/Presentations:

[25] J. Overhold and K. Kearns, "AFRL Autonomy" (presentation, 2013), https://sites.nationalacademies.org/cs/groups/depssite/documents/webpage/deps_084776.pdf

[26] Colin Parris, "Maximizing Value with Industrial AI" (presented at DOE's AI Innovation X Lab, Chicago, IL, October 2019), https://blogs.anl.gov/aixlab/wp-content/uploads/sites/86/2019/10/Maximizing-Value-with-Industrial-AI_10.1.19_Colin-Parris.pdf.

A.6. METHODS FOR CONTINUAL AND ONLINE LEARNING FOR DIGITAL TWINS

Chair:	Nurali Virani	GE Research
Cochair:	Jonathan Ozik	Argonne National Laboratory
Participants:	Panos Stinis	PNNL
	Varun Chandola	University at Buffalo
	Achalesh Pandey	GE Research
	Phil Scruggs	University of Tennessee, Knoxville (UTK)
	Junshan Zhang	Arizona State University

Introduction:

DT models are necessarily limited in scope (not arbitrarily accurate), and data is needed to improve forecasts and update parameters in several scenarios and application domains. Unlike continuous learning in which models are continuously updated with each new data point, continual learning systems identify the need for updating models, obtain suitable data, and then update models. Continual learning is needed to accommodate for concept drifts and input distribution shifts. Although specific approaches for continual learning exist, this is a wide-open research area. This section identifies a few key challenges and some research opportunities for continual and online learning of DTs, which can have a game-changing impact on current and future DT applications for design, forecasting, optimization, and control of various physical systems.

Key challenges:

- In a continual setting, a system must detect input and concept drift (including regime changes, anomaly/novelty detection) to determine if, when, and how to update models.
- The models must be updated by avoiding catastrophic forgetting or enabling intentional forgetting based on system behaviors and environmental effects on the system.
- Techniques are needed to handle diverse timescales for learning, which can change based on system dynamics, data required, and time to update models.
- The system must infer if adequate data is available and if models have learned enough.
- The system must determine where to update models (i.e., at the edge, in the cloud, or in HPC environments) based on resource constraints and data availability.
- The system must determine how to get the correct data from the correct sources (e.g., actual experiments, historical data, virtual experiments, new data from existing streaming sensors, on-demand measurements) to update models.
- Cost and latency of edge devices doing online training and inference must be reduced for scalable adoption, and current learning and inference algorithms are not sufficient.

State of the art:

- Condition-based or performance-based model retraining and incremental learning algorithms are gaining momentum.
- Online learning in edge devices is very limited in scope because current devices primarily support fine-tuning of a few neural network model parameters to update models.
- Reservoir sampling and experience replay techniques are being explored to avoid forgetting.
- Continual learning currently focuses on updating a single model from a single data source. Building flexible multicomponent (i.e., multiple models and multiple data sources) continual learning capabilities with on-demand data gathering is difficult.
- In some industrial settings, feature analysis is done at the edge, and models are executed and updated on a central server with information from extended spatiotemporal horizon.
- Concept drift detection and anomaly detection (AD) approaches that compute the residual between expected and actual measurements are used for performance assessment when actual ground truth information is readily available.

Research opportunities:

- Dynamic model update mechanisms are needed that can identify and execute suitable update mechanisms for fine-tuning, model parameter updates, model architecture adaptation, model feature/sensor set adaptation, or updated data gathering.
- Continual learning of unsupervised models that provide learned representations to track low-dimensional manifolds, which can allow tractable control of complex systems, is needed.
- Developing flexible and generalizable software infrastructure for learning approaches over different sensing and compute platforms is also needed. For example, enabling learning systems with event-triggered workflows for simulation-as-a-service for conducting on-demand ensemble experiments for data generation and updating models will benefit the scientific ML community.
- There is an opportunity to continually update models with decentralized and privacy-preserving collaborative learning across data silos with different levels of access restrictions as new information is available in different silos.
- Although some work exists for incorporating hard or soft constraints in loss terms or in the architecture to include domain constraints, new work is needed to enable continual learning in the presence of domain constraints.
- Learning for stochastic systems by leveraging ideas from nonlinear filtering and signal processing methods is another opportunity.
- In safety-critical and other industrial systems, continual learning must be complemented with continual validation that can characterize error and uncertainty propagation over evolving multicomponent DTs to provide assurance of safety and reliability at run time.

- Efficient methods for learning from spatiotemporal streams are needed.

- Algorithm development for competence monitoring (including online performance assessment) and out-of-distribution detection—when labels or ground truth information is not readily available—is needed to identify the need for updating models.

- Multimodal learning systems that can learn from heterogenous sources of information (e.g., actual experiments, historical data, virtual experiments, new data from existing streaming sensors, on-demand measurements) should be explored.

What will be different if the research opportunities are adequately addressed?

- More reliable, accurate, and competence-aware twins will be available.

- If DTs can stay up-to-date and accurate, they can be used for improved just-in-time planning and monitoring of flexible manufacturing systems (e.g., additive manufacturing [AM]), for which the *new normal* changes frequently.

- Because RL is a key user of continual learning approaches, better continual learning will significantly impact RL. For example, continual learning with domain constraints can help create RL that can satisfy mission, task, and safety constraints.

- More continual learning pipelines will be accessible for scientific and engineering design applications (e.g., edge, HPC, and cloud) in the science and engineering community.

- With reliable adaptive individual models, predictions for multiagent systems can become more accurate.

- Scalable learning for larger networks across boundaries will also be available.

A.7. CONSTRUCTING DIGITAL TWINS USING HIGH-DIMENSIONAL DATA

Chair:	Arvind Mohan	LANL
Cochairs:	Jason St. John	Fermi National Accelerator Laboratory (Fermilab)
	Michael Churchill	Princeton Plasma Physics Laboratory
Participants:	Chris Tennant	Jefferson Laboratory

Introduction:

It is clear that DL can be very successful given quantities of good quality training data and/or physics constraints. But what happens when the data is extremely high-dimensional? This situation frequently occurs in 3D spatiotemporal physics (e.g., turbulence). Such data can often have more than 10^8 degrees of freedom. The few ML efforts that have tried learning such complex spatiotemporal data sets with large degrees of freedom have mostly attempted to apply more compute power. However, efficiency in DL is important for both time-sensitive and cost-sensitive engineering applications; computing resources are neither infinite nor cheap, so a better approach will be needed. How can one develop (1) more efficient neural networks and (2) better ways of representing large data sets so that DL does not require excessive computing resources? Although physics constraints and training data quality have received a lot of attention in DL for large spatiotemporal phenomena, considerably less attention has been paid to the scaling and the efficiency of neural networks. This is not merely an HPC problem, and the community needs fundamental algorithmic advances to truly push DL to embrace realistic engineering problems, however large. These difficulties can be compounded further in DTs when a simulation must be compared with real systems—usually with summary statistics of the two systems. Solutions range from techniques inspired by applied math/information sciences for a parsimonious representation of data, to incorporation of physics constraints in such a way that the resulting neural networks train more readily.

Guiding questions:

- Are there different classes of high-dimensional data?
- Are there efforts and benchmark problems for quantifying neural network efficiency on high-dimensional data?
- What applications are impacted the most by high-dimensional learning and why?
- How is the computational challenge of training overcome when input vectors are so large?
- UQ is always an important issue for which most solutions increase computational needs. It is exacerbated by the computational demand of large data sets. What are some ways to effectively tackle this challenge?
- How can virtual and real systems more effectively make comparisons with high-dimensional data?

Key challenges:

- Training ML models becomes computationally demanding and expensive for several engineering-grade data sets because they are high-dimensional.

- Discussions identified two broad categories of data sets: (1) spatiotemporal physics problems (e.g., turbulent flows, materials physics, and data from radio astronomy) [1] and (2) sensor/instrument data from a large number of heterogeneous sources (e.g., particle accelerators and mega-industrial plants).

- In addition to training neural networks, interpreting models of high-dimensional systems is also challenging. A common strategy is to perform preprocessing with dimensionality reduction. A key challenge there is deciding what information to truncate, discard, or combine.

- How does one enforce system-wide constraints given such data sets? This is important for sensor/plant data in which intrinsic constraints are not readily available in the form of governing equations (unlike several physics problems).

State of the art:

- Current studies approach high-dimensional data as an HPC problem. If more compute capacity can be spared, then the problem will fit. This approach is unfortunately not scalable long term because simulation data (often used as training data sets) are reaching terabyte ranges with exascale computers on the horizon. The computational cost of training on these data sets can be prohibitively expensive in learning reduced-order models, which are often the bedrock of many DT applications.

- On the other hand, strongly enforcing physics constraints in neural networks enable it to learn faster and on lesser data. This indirectly reduces the computational burden. However, explicit efforts to handle high dimensionality have not percolated into mainstream research.

- Popular Sparse 1D-Operational Autoencoders are neural network autoencoders used to learn low-dimensional representations of high-dimensional data sets, often as a preprocessing step. Although they are successful, they paradoxically require significant computational resources [2] to learn these representations. Furthermore, they suffer from typical issues of interpretability and uncertainty, which requires more computational effort to quantify.

- Most simulation/experiment comparisons apply approximate Bayesian computation, which can only be effective with a handful of parameters.

Research opportunities:

- A fruitful direction of research is computing parsimonious representations of data to be fed as input to neural networks, as opposed to the *kitchen-sink* approach of deploying neural networks with tens of millions of parameters directly on a raw data set. This is a problem that borders on applied math and information sciences but is nevertheless of direct significance to the ML community.

- Another important opportunity or goal is to enable a comparison of simulation and experimental data beyond just the summary statistics because this will better enable one to connect phenomena of differing resolutions captured by both techniques. These statistics can assist in reducing the amount of data required for training.

- Processing multisensor data from thousands of sources that do not have well-defined connectivity information (as opposed to spatiotemporal physics) requires a different approach. Graph neural networks are well suited to analyze such data sets and learn connectivity + dependency information that is not apparent to human eyes.

- Despite HPC being the tool of choice for tackling high-dimensional data, there is still room to progress. A fruitful direction would be to quantify the practicality of mixed-precision training for scientific problems, especially in regression-based neural networks. Although much effort has been expended in classification problems, regression is trickier and predominant in science problems.

- Likelihood-free inference using flow-based neural networks for principled comparisons between simulation/experiment is another opportunity.

Impact if research opportunities are addressed:

- Addressing these opportunities can lead to fast and cost-effective DTs for a variety of applications that need rapid modeling (e.g., industrial plants, space applications) under high-risk/consequence scenarios.

- Improved basis identifications and data factorizations can accelerate training and reduce computational cost and add a significant layer of interpretability to high-dimensional data sets.

- Parsimonious data factorizations and lower computational costs would directly benefit UQ of the ML models. This is especially important when using DTs for real-time or rapid decision making, for which multiple ensembles are required to understand uncertainty.

- The rapidly growing field of DTs for earth and climate science applications, in which data is inherently high-dimensional, will be impacted significantly.

Summary:

Although the DT community has rightfully placed significant emphasis on developing ML algorithms that enforce domain-information and physics constraints, there are several constraints in deploying these algorithms to engineering systems of reasonable scale, given cost constraints. ML-based DT development would be better served if tackling high-dimensionality was not considered an afterthought to be handled with HPC but rather as a research problem. Dimensionality reduction has been an active area of research in the physics community for decades, and there are several opportunities blending it with the current crop of physics-based ML research. This is even more important because memory in GPUs—the common workhorse of the ML community—is not as cheap or widely available as CPUs, which contributes to the significant costs of ML training. Finally, there are several applications in earth sciences, astrophysics, and other large-scale problems that are firmly outside the reach of current computational capacity despite considerable interest in ML-based DTs. These domains will benefit significantly from advancements in this area.

Supporting Material:

Papers:

[1] Dayton L. Jones, Kiri Wagstaff, David R. Thompson, Larry D'Addario, Robert Navarro, Chris Mattmann, Walid Majid, Joseph Lazio, Robert Preston, and Umaa Rebbapragada, "Big Data Challenges for Large Radio Arrays," *Proceedings of the 33rd IEEE Aerospace Conference* (March 2012): https://www.wkiri.com/research/papers/jones-bigdata-ieee12.pdf.

[2] Michael Mesarcik, Albert-Jan Boonstra, Christiaan Meijer, Walter Jansen, Elena Ranguelova, Rob V. van Nieuwpoort, "Deep learning assisted data inspection for radio astronomy," *Monthly Notices of the*

Royal Astronomical Society 496, no. 2 (August 2020): pp. 1517–1529, https://doi.org/10.1093/mnras/staa1412.

A.8. CONSTRUCTING AND USING DIGITAL TWINS FOR ANOMALY DETECTION

Chairs:	Adi Hanuka	SLAC National Laboratory
	Jiaxin Zhang	ORNL
Participants:	Mina Sartipi	University of Tennessee at Chattanooga
	Jessica Jones	Sandia
	Hao Huang	GE Global Research

Introduction:

AD is an important problem that has been explored within diverse research areas and application domains. Many AD techniques have been specifically developed for certain application domains, whereas others are more generic. This section provides a structured overview of research on AD. Existing techniques are grouped into different aspects based on the underlying approach adopted by each technique. For each aspect, the group identified key challenges, state-of-the-art technical methods, research opportunities, and impacts, which are used by the techniques to differentiate between normal and anomalous behavior.

Guiding questions:

- Are there standard guidelines for constructing an AD technique that can be applied across different disciplines?
- Are there safety checks recognized by the IoT community to guarantee that AD is self-reliant after deployment?
- Are there common paradigms that every AD must respect?
- How does one reliably incorporate partial knowledge and data gaps in AD construction?
- Are there benchmark tests that are recognized in the AD community as valid pass/fail tests that certify the correct functioning of the AD?

Key challenges:

- High frequency data/miss data, imbalanced data
- Lack of labeled data
- Identifying root cause from upstream data
- Relatively mid-/long-term predictions
- Incorporating uncertainties, reliabilities, and safety
- Improving the robustness of prediction
- Interpreting the anomaly—what makes it an anomaly?

State of the art:

- Promising progress is being made with specific applications using advanced statistical analysis and DL techniques, including supervised and unsupervised methods.
- The supervised approach requires significant effort, and the unsupervised approach relies on assumption and prior information.
- Semi-supervised learning constructs a model that represents normal behavior from normal data and then tests the likelihood of a test instance.

Research opportunities:

- Data preprocessing can leverage recent advances (e.g., dimension reduction) in ML/AI to handle multiple data types, including time-series data, image-based data (e.g., computed tomography [CT]), and video data.
- A scientific AD benchmark data set (experimental or simulation approach) for supervised (convolutional neural networks, RNN/LSTM, graph learning) and unsupervised (clustering analysis, generative adversarial network [GAN], normalizing flow) approaches is one research opportunity.
- An opportunity exists to incorporate prior/domain knowledge in a rigorous way and improve the robustness given sparse data.

Impact if research opportunities are addressed:

- Enable the investigation of super large data sets with high frequency and high resolution.
- Improve the prediction accuracy, efficiency, and robustness with higher confidence, specifically when given a sparse labeled data set.
- Enable a fair comparison and evaluation of new AI/ML methods using a scientific benchmark data set instead of MNIST or imagenet.
- Potentially extend to other scientific applications and related domains (e.g., adversarial attacks).

Supporting Material:

Papers:

[1] Chris Tennant, Adam Carpenter, Tom Powers, Anna Shabalina Solopova, Lasitha Vidyaratne, and Khan Iftekharuddin, "Superconducting radio-frequency cavity fault classification using machine learning at Jefferson Laboratory," in *Phys. Rev. Accel. Beams* 23, no. 11 (November 2020): https://doi.org/10.1103/PhysRevAccelBeams.23.114601.

[2] E. Fol, R. Tomás, J. Coello de Portugal, and G. Franchetti, "Detection of faulty beam position monitors using unsupervised learning," in *Phys. Rev. Accel. Beams* 23, no. 10 (October 2020): https://doi.org/10.1103/PhysRevAccelBeams.23.102805.

[3] Kevin M. Potter, Brendan Donohoe, Benjamin Greene, Abigail Pribisova, and Emily Donahue, "Automatic detection of defects in high reliability as-built parts using x-ray CT," in *Proceedings*

Volume 11511, Applications of Machine Learning 2020, 115110O (2020): https://doi.org/10.1117/12.2570459.

[4] Michael C. Krygier, Tyler LaBonte, Carianne Martinez, Chance Norris, Krish Sharma, Lincoln N. Collins, Partha P. Mukherjee, and Scott A. Roberts, "Quantifying the unknown impact of segmentation uncertainty on image-based simulations," arXiv:2012.09913, https://arxiv.org/abs/2012.09913.

A.9. HARDWARE AND SOFTWARE ISSUES IN EDGE COMPUTING AND PRODUCTION DEPLOYMENT OF DIGITAL TWINS

Chair:	Ron Oldfield	Sandia
Cochairs:	Iris Bahar	Brown University
	Mike Lang	LANL

Introduction:

Computing/sensing at the edge is an important aspect of supporting DT capabilities, and the DOE laboratories and its partners could (and should) play an important role in evolving edge-computing technologies. DTs provide a virtual representation of real-world objects or environments, and it is critical that the real-world/physical objects interact with and inform the virtual representations, particularly in resource-constrained environments in which the edge systems may be limited by power, weight, computing capabilities, and response time, and the communication capabilities between the digital and physical twin might be limited by communication. Examples like this exist in DTs of experimental facilities, remote sensing (e.g., satellites), autonomous vehicles, and many others.

Key challenges:

The team identified distinctly different challenges for edge deployment depending on the DT's phase of development.

- Edge computing in DT design phase: During the design of the digital and physical twin, a key role for edge computing is to validate and provide confidence in the virtual representation of the device and inform the designers about requirements and expectations for deployed edge devices. For example, a DT of an autonomous vehicle may require several well-placed sensors to understand the impact of environmental conditions on internal systems (e.g., electronics). The DT could help determine where to place sensors, what data must be processed at the edge, and what data must be communicated to the DT.

- Edge computing for deployment: Once in the field, requirements and constraints for edge computing are different than during the design phase. In particular, limitations in power/energy, communication, and the ability to reconfigure the system are expected, and security/privacy are much more important.

State of the art:

- Edge Hardware
  - The edge device (i.e., hardware) typically refers to the processing element closest to the edge sensor (or in some cases on the sensor). Requirements for this device depend on the type and volume of data captured by the sensor, the power available for processing, and communication and connectivity requirements between the edge device and the DT. Edge hardware is typically lightweight and custom. The current state of the art for the edge device design leverages devices like Raspberry Pi and ARM processors for prototypes and relies on custom architectures for deployment to meet design constraints.

- Edge Software
  - Depending on the application, edge devices perform signal processing, data fusion, data reduction, data analytics/prediction, noise reduction, data labeling, and much more. In most cases, the edge device must be able to process, manage, and communicate results to the DT. Secure

computation and communication may also be required. The IoT vendors are evolving development environments for edge devices. Some examples include Android Things (Google), Predix (GE), and Azure IoT Suite (Microsoft).

Research opportunities:

- Hardware/sensor design
  - Opportunities exist to expand DOE investments in power-efficient, secure, and communication-efficient hardware (e.g., quantum information science, neuromorphic hardware, reconfigurable computing, system-on-a-chip).

- Robustness, security, and resilience
  - Opportunities exist to assess vulnerability and fragility of neural-network models deployed to edge devices.
  - Adapting to and recovering from failure, data corruption, and adversarial attacks.
  - Employing and developing algorithms for secure and robust communication and computations among networks of sensors (fog computing).

- Autonomy
  - Hardware and algorithms for rapid response and in situ decision making on edge devices and systems.

- Adapting to changing requirements of the DT
  - Development of agile and reconfigurable edge devices and software. How do you update hardware and software on a remote edge device (e.g., a satellite)?

Impact if research opportunities are addressed:

- Better understanding of hardware/software requirements for edge devices in support of DT.

- Custom edge hardware designs that satisfy resource constraints and provide sufficient capability to support DTs.

- Evolving approaches for improved robustness, security, and resilience in edge systems.

- DT workflows that enable continuous updates and information exchange with the edge system.

- Improvements to ML algorithms and continual learning approaches.

Summary:

Edge systems play an important role in the deployment of DTs. The DOE laboratories as well as industry and academic partners could play an important role in extending the state of the art through foundational research hardware and software for resource-constrained edge devices. Such efforts will improve the ability to leverage DT technology to address key issues in DOE's science, energy, and national security missions.

Supporting Material:

Papers:

[1] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge Computing: Vision and Challenges," in *IEEE Internet of Things Journal* 3, no. 5 (October 2016): pp. 637–646, http://doi.org/10.1109/JIOT.2016.2579198.

[2] Y. He, J. Guo, and X. Zheng, "From Surveillance to Digital Twin: Challenges and Recent Advances of Signal Processing for Industrial Internet of Things," in *IEEE Signal Processing Magazine* 35, no. 5 (September 2018): pp. 120–129, http://doi.org/10.1109/MSP.2018.2842228.

[3] M. Capra, R. Peloso, G. Masera, M. Ruo Roch, and M. Martina, "Edge Computing: A Survey on the Hardware Requirements in the Internet of Things World," *Future Internet* 11, no. 100 (2019), https://doi.org/10.3390/fi11040100.

A.10. SCALING AND REALIZATION OF DIGITAL TWINS ON CLOUD-AND-HPC SYSTEMS

Chair:	Piyush Modi	NVIDIA
Cochair:	Sudip K. Seal	ORNL
Participants:	Sylvain Bernard	Sandia
	Nathan DeBardeleben	LANL
	Hal Finkel	DOE
	Jason Hick	LANL
	Ron Oldfield	Sandia
	Srikanth Yoginath	ORNL

Introduction:

DTs of complex systems involve millions of industrial systems (e.g., factories, fleets of transport systems, power plants), billions of assets and an even larger number of subcomponents (e.g., valves, pumps) with which they are integrated. Hence, the ability to scale the training and inferencing workloads to build such DTs requires highly scalable compute infrastructure that facilitates edge processing of sensor data to cloud and HPC data center processing of simulation, model finding, training, inferencing, and continuous learning. To facilitate such workloads, orchestration stacks must evolve from traditional HPC schedulers to service-oriented DL/ML orchestrators based on Kubernetes (K8s) to a hyperconverged stack that facilitates the construction of DT and the deployment and related lifecycle management of DT models.

Guiding questions:

DT Definition

- What roles do DTs play for DOE facilities, multicomponent engineered systems (e.g., detectors, accelerators), emerging technologies (e.g., federated instruments), and novel workflows (e.g., edge computing)?

- What does scale mean in the context of DT? Is it the number of twins, training cycles, update cycles, simulation fidelity, simulation latency, number of users, infrastructure, number of systems, or number of subsystems?

- Is a DT needed for every salient component or as one catch-all DT?

- What is the data strategy for scalable DT?

- How does one identify the DT compute pipeline, infrastructure, and workflow?

- What are the appropriate development operations and ML operations in the context of a DT model lifecycle (training, deployment, and continuous updates)?

DT Scale

- Data
  - How does one identify all the data sets and their sources?
  - Where do they reside (HPC vs. big data infrastructure vs. AI training data)?

    - How does one virtualize access for different users based on their needs?
    - Can a common HPC-cloud data-management infrastructure to support ingest/correlation of experimental and mod-sim data be established?
    - What are the best practices and processes for data ingestion and access from the cloud?

- Workflows and Workloads
    - For DT workflows, how does one deploy multiple GPUs on single node and multinode workflows (potentially with multiple GPUs on each node) that are reminiscent of OpenMP/threads on the node and MPI/PVM off the node?
    - Can smart DT workflows be deployed that can train one model per GPU per core or share data between GPUs/cores or even across nodes?
    - How does one expose HPC systems to look like a cloud service?
    - How does one orchestrate an ensemble of models?
    - How does one load balance between them while maintaining efficiency?

Key challenges:

DT Definition

- What are the relevant assets or processes to model?
- What are the representation technologies to compose/orchestrate system of systems models?
- What are the relevant objectives (AD, classification, prediction, prescription, surveillance, and process optimization)?

Compute

- How does one evolve HPC-optimized compute infrastructure and tools for AI, DL, and DT?
- Does the scale and complexity of DT directly correspond to data required or data to be streamed?
- How does one support HPC systems with high-bandwidth external networks (Internet2) for streaming data from geographically separated physical sites and DTs? This is compulsory for earth system processes.

Data

- How does one determine the relevance of data and how often to collect?
- How does one build a scalable and fault-tolerant data storage infrastructure?

Scale

- How does one scale ML algorithms for training and updating DTs?
- How does one carry out large-scale training on streaming data and/or historical data?
- How does one execute large-scale simulations with ML model support to achieve higher accuracy?

- How does one construct anticipatory DT—a computational framework for AI to perform additional simulation runs based on the expectation of a certain high-impact event occurring in the future—and be able to answer/address questions for unfolding behavior caused by the high-impact event way ahead in time?

- How does one deploy RL frameworks to train AI agents to perform necessary control decisions under low false–positive and no false–negative conditions?

- How does one build interconnected dynamic systems able to predict single or cascading failures?

- How does one build AI pipelines for varied DT use cases and related data and performance challenges, including the following?

    - Automation for data curation, augmentation, and feature engineering
    - AD
    - Future prediction (e.g., remaining useful life)
    - Scenario planning and pattern search
    - Sparse data, missing data, and UQ
    - Fidelity and latency

State of the art:

- DL-related benchmarks are tracked for accuracy and performance by MLPerf.org.

- AD solutions using autoencoders and other generative techniques are being adopted, and the Massachusetts Institute of Technology's (MIT's) TADGAN is the most recent.

- Failure prediction recurrent networks ranging from LSTM and GANs to transformers are also being adopted.

- Mining domain-specific entities, summarization, and context-aware knowledge from textual data in the industrial environment, BERT and related NLP tools are being adopted for domain-specific language modeling and are being fine-tuned for various tasks.

- Frameworks:
    - Tensorflow, PyTorch, MxNet, and RAPIDS.ai are some of the software frameworks.
    - Horovod is a meta-level platform to scale training across GPUs and nodes.
    - NVIDIA's SimNet is evolving to facilitate physics-driven simulations augmented by neural networks to train physics-influenced neural networks.
    - NVIDIA Omniverse is a workbench used to assimilate data, design DT assets in simulation, facilitate collaboration, and facilitate deployment of DTs.
    - NVLink scales interGPU networking up to 600 Gbps.
    - Melanox Infiniband is an end-to-end high-speed Ethernet and InfiniBand interconnect solution.
    - Mellanox Rivermax uses GPUDirect for optimal data flow from network to GPU memory.

- Hyperconvergence facilitating stacks have evolved to orchestrate DT workloads from training; inferencing; and active, federated, and continuous learning (e.g., Google Anthos, Amazon Outpost, Azure Stack, VMWare Tanzu, NVIDIA GPU Operator, K8, Helm Charts, NVIDIA NGC, SLURM/Singularity integration via K8 Project – WLM-Operator and related machine learning operations stacks).

- Compute technology readiness and availability
  - GCP, Azure, Amazon Clouds
  - ORNL's Summit
  - NVIDIA DGX Superpod.

Research opportunities:

- Development of AI workbench tools to enable data ingestion, curation, collaboration, and visualization.
- Development of HPC hardware and software stacks and computing infrastructure specifically geared toward deployments of DT.
- Development of efficient scalable workflows.
- Opportunities to combine and approximate physics with data-driven techniques.
- Opportunities for transfer and federated learning for science.
- Development of advanced workflows for user engagement/usability.
- Advancements in scalable DL-aided and physics-based simulations in which DL models are either used to accelerate or improve model accuracy of physics-based simulations.
- Opportunities to leverage HPC resources to meet real-time operational requirements.
- Opportunities to leverage HPC resources to realize DT for DOE facilities (e.g., ORNL's Spallation Neutron Source [SNS]) that generate large amounts of operational-specific data.
- Advancing automated data annotation and labeling techniques and interpretable features.
- Study of AI assurance, UQ, fraud prevention, and detection.

Impact if research opportunities are addressed:

- DTs and physical twins can automatically update and exchange information at scale (i.e., train, update, and deploy models automatically).
- 3D visualization and interaction with physical and DT systems.
- Ability to proactively learn *what-if* scenarios in anticipation of possible future events.
- HPC-aided DT realization for large DOE facilities.
- DT construction tools for scientists that facilitate multidisciplinary collaboration.

Summary:

Development and deployment of large-scale DL-based DT on cloud and HPC systems entails the aggregation of data sets residing in siloed infrastructure and the orchestration and streamlining of

compute infrastructure managed by modern ML and development operations for execution of a wide range of legacy and emerging algorithms. Facilitating continuous learning to maintain DT performance and to accurately mirror the DT's physical counterpart is an overarching goal of DT development and deployment. To attain these goals, the community must have AI workbench tools that can ingest data from legacy systems with varied formats and can universally represent their semantics in the context of the physical world and associated digital world tags to trigger the execution of DT workflows. Several state-of-the-art tools, algorithms, and related benchmarks have been identified. They also represent related research opportunities for software tools, hyperconverged compute resources, storage and networking infrastructure, AI/DL/ML/PINN algorithms, and interactive user interfaces. Realizations of such scalable DTs will require the community to adopt new tools, evolve their traditional HPC workflows by adopting comprehensive DT design, and develop and deploy processes with emerging hyperconverged edge-to-cloud compute infrastructure to execute AI/DL/PINN algorithms at scale.

DT conception, development, and deployment should be based on realistic operational needs of high-end scientific instruments and large-scale facilities hosted by the DOE. DTs can play a significant role in constantly surveilling scientific systems to track operational deviations, predict early failures, detect cyber intrusions, and to devise process optimizations. Facilities such as SNS, the High Flux Isotope Reactor, particle accelerators, and several high-end sophisticated scientific instruments can benefit from a DT-asserted operational consistency. There are commercial off-the-shelf products that aid in DT development and deployment, but the type of technology to use for realizing DTs for large-scale DOE scientific instruments is also dependent on the system's operational security level.

Supporting Material:

Papers:

[1] Yongyi Ran, Xin Zhou, Pengfeng Lin, Yonggang Wen, Ruilong Deng, "A Survey of Predictive Maintenance: Systems, Purposes and Approaches," arXiv:1912.07383v1, https://doi.org/10.48550/arXiv.1912.07383.

[2] Annual Conference of the Prognostics and Health Management Society 2020, virtual, November 2020, https://phmsociety.org/conference/annual-conference-of-the-phm-society/annual-conference-of-the-prognostics-and-health-management-society-2020/.

[3] Alexander Geiger, Dongyu Liu, Sarah Alnegheimish, Alfredo Cuesta-Infante, and Kalyan Veeramachaneni, "TadGAN: Time Series Anomaly Detection Using Generative Adversarial Networks," arXiv:2009.07769, https://doi.org/10.48550/arXiv.2009.07769.

[4] Bryan Lim, Sercan O. Arik, Nicolas Loeff, and Tomas Pfister, "Temporal Fusion Transformers for Interpretable Multi-horizon Time Series Forecasting," arXiv:1912.09363, https://doi.org/10.48550/arXiv.1912.09363.

[5] S. Yoginath et al., "On the Effectiveness of Recurrent Neural Networks for Live Modeling of Cyber-Physical Systems" (2019 IEEE International Conference on Industrial Internet [ICII], Orlando FL, November 2019), https://doi.org/10.1109/ICII.2019.00062.

[6] NVIDIA Corporation, *NVIDIA DGX SuperPOD: Scalable Infrastructure for AI Leadership*, RA-09950-001, October 2021, https://images.nvidia.com/aem-dam/Solutions/Data-Center/gated-resources/nvidia-dgx-superpod-a100.pdf.

[7] K. Perumalla, S. Yoginath, and J. Lopez, "Detecting Sensors and Inferring their Relations at Level-0 in Industrial Cyber-Physical Systems" (2019 IEEE International Symposium on Technologies for Homeland Security [HST], Woburn, MA, March 2019), https://doi.org/10.1109/HST47167.2019.9032891.

Blog Posts:

[8] Ian Lumb, "Introducing HPC Affinities to the Enterprise: A New Open Source Project Integrates Singularity and Slurm via Kubernetes," *Syslabs.io*, May 7, 2019, https://medium.com/sylabs/introducing-hpc-affinities-to-the-enterprise-a-new-open-source-project-integrates-singularity-and-6461091c2626.

[9] Colin Parris, "Digital Twin 2.0 and the emergence of 'Humble AI,'" *LinkedIn*, January 16, 2019, https://www.linkedin.com/pulse/digital-twin-20-emergence-humble-ai-colin-parris/.

[10] Amy Kover, "Humble AI Takes A Curious Turn: How Algorithms That Ask 'Why' Can Improve Wind Energy," *GE*, November 18, 2019, https://www.ge.com/news/reports/humble-ai-takes-a-curious-turn-how-algorithms-that-ask-why-can-improve-wind-energy.

[11] Nefi Alarcon, "Accelerating Automated and Explainable Machine Learning with RAPIDS and NVIDIA GPUs," *NVIDIA Technical Blog*, November 17, 2020, https://developer.nvidia.com/blog/accelerating-automated-and-explainable-machine-learning-with-rapids/.

Software Repositories:

[12] "RAPIDS Notebooks," GitHub repository, last commit February 4, 2022, https://github.com/rapidsai/notebooks.

Websites:

[13] "Time Series Analysis," RAPIDS AI, https://medium.com/rapids-ai/tagged/time-series-analysis.

[14] "NVIDIA TAO Toolkit," NVIDIA, https://developer.nvidia.com/tao-toolkit.

[15] "NVIDIA Modulus: A Framework for Developing Physics Machine Learning Neural Network Models," NVIDIA, https://developer.nvidia.com/modulus.

[16] "Machine learning innovation to benefit everyone," ML Commons, updated December 14, 2021, https://mlcommons.org/en/.

[17] "RAPIDS: Open GPU Data Science," RAPIDS, updated February 2022, https://rapids.ai/.

[18] "Develop with NVIDIA Omniverse," NVIDIA, https://developer.nvidia.com/nvidia-omniverse-platform.

[19] "NVIDIA ISAAC: The Accelerated Platform for Robotics and AI," NVIDIA, https://www.nvidia.com/en-us/deep-learning-ai/industries/robotics/.

A.11. NUCLEAR ENERGY: CHALLENGES AND APPLICATIONS OF DIGITAL TWINS

Chair:	Prashant K. Jain	ORNL
Cochair:	Vaibhav Yadav	Idaho National Laboratory
Participants:	Bob Ledoux	US DOE, ARPA-E Program
	Raj Iyengar	US Nuclear Regulatory Commission
	Doug Eskins	US Nuclear Regulatory Commission
	Pradeep Ramuhalli	ORNL
	Vittorio Badalassi	ORNL
	Justin Weinmeister	ORNL

Introduction:

The purpose of this breakout session was to better understand the potential applications of DTs in nuclear power plant operations, identify the associated technical challenges, determine potential solutions, and assess the regulatory viability. In this digital and data-driven age, it is conceivable that nuclear plants could operate autonomously with limited or remote human interference. The primary enablers of such futuristic plants are

- the advanced and comprehensive sensor network that would generate operational and system data and provide a data management module to manage the data,
- a physics-based analytical backbone to analyze relevant data and render predictions, and
- an AI or ML engine to parse through the analyses predictions to provide recommendations and take necessary actions.

Guiding questions:

- What values are expected out of DTs for a nuclear system? Where would DTs provide maximum benefit in the life cycle of a nuclear plant? What applications (both safety-significant and not safety-significant) within a nuclear power plant (e.g., design, construction, operation, maintenance, and decommissioning) would benefit most from a DT? [The Future Value Proposition]
- What are some major challenges in building and deploying DTs for nuclear systems? What are some unique needs of a compliant DT technology for nuclear applications? Are there any specific V&V issues associated with deploying/using DT in a regulated environment? Which of these are specific to ML for DT? How might these issues be addressed given the current state of the art in ML? [Key Challenges]
- What are major elements within a nuclear power plant (light water and advanced reactors) for which developing a DT might make sense? What would be a low-hanging adoption approach for DTs in nuclear applications? What existing digital nuclear support systems can be evolved into standing up a functional DT? [State of the art]
- Given that this is an ML for DT workshop, what are the various ways in which ML might be injected into the DT for nuclear applications? What bridging technologies must be developed? What should be the near-term (5 years) to long-term (>10 years) research focus? [Research Opportunities]

- Data is a key need for ML. What gaps exist concerning data needed for a robust DT for nuclear applications? How might these gaps be addressed? For example, what are the simulation, test beds, instrumented plants, sources, and data management infrastructure needs for managing data sets? Is there a strategy to collect and build data repositories from ongoing experimental programs? [Research Opportunities]

Key challenges:

- Interoperability (composability, scalability, heterogeneity): bridging of physics and data-driven ML/AI models for a range of coupled subsystems) [building twins]

- Reliability (accuracy, predictability, assurance, explainability, interpretability) [qualifying twins]

- Sustainability (maintainability, adaptability, robustness, reconfigurability, learnability) [deploying twins]

- Security (cyber and information security, confidentiality, operational controls, safety implications of data breaches and hacks) [securing twins]

State of the art:

- DTs are gaining broader attention in the nuclear space through an emerging interest in digital technologies within the advanced reactors industry and with research support from the DOE ARPA-E Generating Electricity Managed by Intelligent Nuclear Assets program.

- However, at present, there is a heavy reliance on empirically (experimental) observed correlations for design, safety, and controls purposes, with leanings toward multiphysics high-fidelity coupled numerical simulations.

- Data-driven hybrid approaches will be a paradigm shift for the nuclear industry and will demand considerable applied R&D investments before they can be successfully adopted.

Research opportunities:

- Generation of synthetic and operational DT-grade data sets and methods to qualify DTs and establishing best practices for long-term archival DT management.

- Development and qualification of physics-informed and data-driven surrogate models and emulators for nuclear design and safety with hybrid co-simulators linked to hardware in the loop.

- Assessing the reliability of DT technologies (e.g., identifying, predicting, quantifying uncertainties), and analyzing the novel failure modes that DTs could potentially introduce.

- Demonstration of early adoption on nonsafety-grade subscale systems (e.g., the power conversion system or feedwater control system).

- Identify a graded approach to expand the scope of DT demonstrations beyond non-safety systems for nuclear assets.

- Developing tools to visualize and make explicit relationships and dependencies among subsystems to provide a holistic operational viewpoint.

Impact if research opportunities are addressed:

- Enhanced adoption and reliance on digital technologies within the nuclear industry; DTs could become a reliable training resource for the workforce (from education to operation to decommissioning).

- A well-designed building information modeling–supported DT model can significantly reduce uncertainties in upfront construction cost and lead times for deployment of a new nuclear facility and its operations and maintenance.

- Real-time remote monitoring and control and autonomous operation of nuclear (micro)reactors would be realized.

- Using the DT as a novel source of regulatory information would increase safety while decreasing licensee regulatory burden (e.g., integrated information, more targeted regulatory activities, fewer on-site inspections, fewer ancillary information requests).

Summary:

The emerging DT technology could support existing and future nuclear reactors and can play a significant role in reducing overall costs while improving operational safety and performance. However, there exist several challenges for the DT concepts in nuclear systems. One major challenge is the complexity and qualification of the meta-models (e.g., hybrid data-driven feedback and decisions) that are employed within DTs. Besides, the long-term viability of a DT is a must have for nuclear systems because of their long-life expectancy (40+ years). Therefore, DTs for nuclear applications must be designed for long-term maintenance of associated computing hardware, software, knowledge management, and cybersecurity framework. Interoperability and scalability of a DT is another major issue because DT combines multiple modeling and simulation tools (physics informed with data-driven insights) with varying degrees of pedigree. However, most of these challenges can be resolved through further technological advancements and research and with early adoption and demonstration successes.

A.12. DIGITAL TWIN CERTIFIED ADDITIVE MANUFACTURING

Chair:	Aric Hagberg	LANL
Cochair:	Luke Scime	ORNL
Participants:	Garrison Flynn	LANL
	Mike Grieves	Florida Institute of Technology
	Craig Miller	Ansys, Inc.

Introduction:

AM (additive manufacturing), or 3D printing, is a promising new manufacturing process that builds parts layer-by-layer. Nominally, AM is ideal for fabricating complex geometries in low to medium production volumes for industries such as aerospace, biomedical, energy production, and automotive. However, the high variability observed in the AM process, particularly in microstructure and defects, is slowing broader adoption within the manufacturing community.

Because fully implemented DTs would allow each unique AM component to be modeled, the quality of each part could be estimated based on specific, as opposed to aggregated, data. Such an approach would allow individual parts to be qualified for safety-critical applications even if the AM process remains significantly variable. Additionally, visualizations and simulations that leverage the DT concept may go a long way toward establishing trust in the AM process and overcoming the cultural inertia present in many of the relevant industries.

In addition to AM benefiting substantially from DTs, AM is an excellent application for the development, testing, and validation of new DT technologies. In particular, the layer-wise nature of the AM process allows for the collection of vast sums of in situ processing data that can be incorporated into a DT. Furthermore, the AM simulation community is extremely active (focusing on thermal, thermomechanical, and microstructural evolution modeling) and will be able to immediately leverage any new DT advancements.

Guiding questions:

- How can one develop confidence that a given AM component will achieve the required performance levels for its target application, including safety-critical applications?

- Can sufficient and appropriate data from an AM build be identified and collected such that virtual modeling and simulation capabilities can be developed and used in place of physical testing?

Key challenges:

- AM processes experience a high degree of interpart and intrapart variability—particularly with respect to defect populations and microstructural features.

- Current AM inspection processes are typically performed post-build (e.g., x-ray CT) and are extremely time consuming and cost prohibitive for most applications. Virtual analysis and testing performed on a DT model would enable increased acceptance of AM components.

- Significant cultural inertia exists in many industries, which must first be overcome before AM can be used to its full potential.

State of the art:

- Post-build part qualification and certification is a major expense and impediment to broader AM adoption.
- Currently, the manufacturing industry only collects fairly coarse-grained data during part fabrication (e.g., using traveler forms).
- Although AI can effectively analyze the massive quantities of in situ data collected during the process, it can be very expensive to collect ground truths for the training data (e.g., tensile tests).

Research opportunities:

- Adding sensors to AM machines to collect in situ process data that can be stored in a DT and analyzed using AI for defect detection and property prediction.
- Enabling virtual, application-specific test to destruction by modeling the DTs and building a knowledge base that validates the virtual test against the physical test.
- Leveraging DTs and AI for real-time process control and defect healing.
- Research into material properties, material genomics, and material development.

Impact if research opportunities are addressed:

- Enable better design for AM to help engineers design better parts that are optimized for AM processing.
- Manufacture parts that are born-qualified in small-to-medium volumes for safety-critical applications in regulation-intensive industries.
- Enable most of what AM promises.

Summary:

The current methodology of physical testing and inspection is cost prohibitive for AM components—particularly components that must be safety rated and produced in small quantities. DTs have the potential to address this issue by moving testing from the physical environment to the virtual environment. The least expensive and least risky inspection and testing capability is to have the information of a component that has been tested to destruction. Unfortunately, collecting that information renders that component unusable. If the specific DT that captures the required data from the AM-built component can be tested to destruction with the same results as testing the physical component (DT Certified), then there can be a high level of confidence as to the future performance of that component. This capability will require research and development of both in situ sensing, modeling and simulation capabilities, and a cultural shift in acceptance of novel testing methodologies.

Supporting Material:

Papers:

[1] T. DebRoy, W. Zhang, J. Turner, and S.S. Babu, "Building digital twins of 3D printing machines," *Scripta Materialia* 135 (2017): pp. 119–124, https://doi.org/10.1016/j.scriptamat.2016.12.005.

[2] Robert X. Gao, Lihui Wang, Moneer Helu, and Roberto Teti, "Big data analytics for smart factories of the future," *CIRP Annals* 69, no. 2, (2020): pp. 668–692, https://doi.org/10.1016/j.cirp.2020.05.002.

[3] D. Mies, W. Marsden, and S. Warde, "Overview of Additive Manufacturing Informatics: 'A Digital Thread,'" *Integr Mater Manuf Innov* 5, (2016): pp. 114–142, https://doi.org/10.1186/s40192-016-0050-7.

APPENDIX B. AIRES 2 WORKSHOP PROGRAM

Artificial Intelligence for Robust Engineering and Science

AIRES 2: Machine Learning for Robust Digital Twins

January 19–21, 2021

Program Committee

General Chair: David Womble, Oak Ridge National Laboratory (ORNL)
Logistics and Planning Chair: Christy Hembree, ORNL

Iris Bahar, Brown University
Kevin Cao, Arizona State University
Frank Liu, ORNL
Dan Lu, ORNL
Justin Newcomer, Sandia National Laboratories (Sandia)
Laura Pullum, ORNL
Pradeep Ramuhalli, ORNL
Abhinav Saxena, GE Research
Malachi Schram, Pacific Northwest National Laboratory (PNNL)
Sudip Seal, ORNL
Dali Wang, ORNL

Robust Engineering is the process of designing, building, and controlling systems to avoid or mitigate failures. The introductory Artificial Intelligence for Robust Engineering and Science (AIRES) workshop in January 2020 explored these foundations. This second AIRES workshop will build on the success of the first workshop to explore and develop the foundations of AI for constructing, deploying, and assuring the robustness of digital twins (DTs). The workshop is composed of three tracks.

AGENDA

Tuesday, January 19, 2021

11:00–11:15 a.m.	Welcome, Introduction, and Agenda Overview Jeff Nichols, ORNL David Womble, ORNL
11:15–12:00 p.m.	Keynote Speaker: Michael Grieves, Florida Institute of Technology *Intelligent Digital Twins: The Role of AI and ML in the Future of Digital Twins*
12:00–12:30 p.m.	Break

Track 1: Construction of DTs

This track will explore the mathematical and computational aspects of using machine learning (ML) to construct robust models of physical systems with an emphasis on dynamical and complex systems. Topics of interest include but are not limited to the following:

- Feature engineering and knowledge representation
- Integrating time-series data for anomaly detection (AD) and trends predictions
- Incorporating physics-based prior information
- Developing an evolving DT through continuous learning
- Data management

Session Chair: Justin Newcomer, Sandia
Session Cochair: Malachi Schram, PNNL

12:30–1:00 p.m.	Invited Speaker 1–1: Nathan Kutz, University of Washington *Targeted use of deep learning for physics and engineering*
1:00–1:30 p.m.	Invited Speaker 1–2: Farinaz Koushanfar, University of California San Diego *Robust and private machine learning*
1:30–1:45 p.m.	Speaker 1–1: Eric Darve, Stanford University *Machine learning for inverse modeling in mechanics*
1:45–2:00 p.m.	Speaker 1–2: Luke Scime, ORNL *Creating scalable digital twins for advanced manufacturing*
2:00–2:15 p.m.	Speaker 1–3: Rose Yu, University of California San Diego *Physics-guided AI for learning spatiotemporal dynamics*
2:15–2:30 p.m.	Speaker 1–4: WaiChing Sun, Columbia University *Microstructure-sensitivity plasticity inferred via graph neural network*
2:30–3:00 p.m.	Break

Session Chair: Pradeep Ramuhalli, ORNL
Session Cochair: Iris Bahar, Brown University

3:00–3:30 p.m.	Invited Speaker 1–3: George Em Karniadakis, Brown University *DeepM&Mnet: A new neural network architecture based on operator regression for digital twins*
3:30–3:45 p.m.	Speaker 1–5: Piyush Modi, NVIDIA Corporation *Tools to accelerate design, development, and deployment of digital twins*
3:45–4:00 p.m.	Speaker 1–6: David Schmidt, University of Massachusetts Amherst *Accelerated DL representation of turbulent, reacting flow*
4:00–4:15 p.m.	Speaker 1–7: Jeph Wang, Los Alamos National Laboratory (LANL) *Digital twins for x-ray and neutron cameras*
4:15–4:55 p.m.	Breakout Sessions
4:55–5:00 p.m.	Day 1 Wrap-Up

Wednesday, January 20, 2021

11:00–11:15 a.m.	Day 2 Welcome and Introduction

Track 2: Application and Deployment of DTs

This track focuses on the practical challenges when using DTs:

- Edge deployment for real-time and power-efficient deployment of DTs
- Federated learning for privacy or for data reduction
- Integrating high-performance computing (HPC) and edge systems, including model and data management
- Online and offline continuous learning on edge-based systems
- Human-machine interface design

Session Chair: Kevin Cao, Arizona State University
Session Cochair: Dali Wang, ORNL

11:15–11:45 a.m.	Invited Speaker 2–4: Felipe Viana, University of Central Florida (UCF) *Digital twins for prognosis applications with hybrid physics-informed neural networks*
11:45–12:00 p.m.	Speaker 2–8: David Stracuzzi, Sandia *Preliminary work on a digital twin for cancer patients*
12:00–12:15 p.m.	Speaker 2–9: Chetan Kulkarni, KBR, Inc., NASA Ames Research Center *Hybrid model-based approaches for systems health management and prognostics*
12:15–12:30 p.m.	Speaker 2–10: Sandra Biedron, University of New Mexico and Element Aero *Experiences in dynamic systems—how we better model, understand, and control intelligently*
12:30–1:00 p.m.	Break

Session Chair: Abhinav Saxena, GE Research
Session Cochair: Malachi Schram, PNNL

1:00–1:30 p.m.	Invited Speaker 2–5: Draguna Vrabie, PNNL *Deep learning digital twins for model predictive control*

1:30–2:00 p.m.	Invited Speaker 2–6: Junshan Zhang, Arizona State University *Edge intelligence in IoT ecosystems: From continual learning to collaborative learning*
2:00–2:15 p.m.	Speaker 2–11: Hao Huang, GE Research *Industrial data anomaly detection and diagnosis with variable association change*
2:15–2:30 p.m.	Speaker 2–12: Jibonananda Sanyal, ORNL *Transportation/mobility digital twin for Chattanooga*
2:30–2:45 p.m.	Speaker 2–13: Jason St. John, Fermi National Accelerator Laboratory (Fermilab) *Digital twins for the Fermilab particle accelerator complex*
2:45–3:00 p.m.	Speaker 2–14: Abha Moitra, GE Research *Automating construction of formal assurance case fragments*
3:00–3:30 p.m.	Break
3:30–4:55 p.m.	Breakout Sessions
4:55–5:00 p.m.	Day 2 Wrap-Up

Thursday, January 21, 2021

11:00–11:15 a.m.	Day 3 Welcome and Introduction

Track 3: Techniques to Provide Assurance

This track will address issues of assuring appropriately designed, constructed, and deployed DTs with a level of rigor consistent with the intended use, including the level of risk. Assurance should be broadly interpreted to include the following:

- Verification, validation, and calibration
- Security and resilience
- Uncertainty quantification (UQ)
- Causal inference
- Detecting and dealing with bias
- Explainability and interpretability

Session Chair: Dan Lu, ORNL
Session Cochair: Sudip Seal, ORNL

11:15–11:45 a.m.	Invited Speaker 3–7: Nurali Virani, GE Research *Humble AI for competence-aware digital twins*
11:45–12:00 p.m.	Speaker 3–15: Jaideep Ray, Sandia *Assembling training data sets for generalizable machine-learned models of physical phenomena*
12:00–12:15 p.m.	Speaker 3–16: Varun Chandola, University at Buffalo *Anomaly detection and clustering for evolving data streams*
12:15–12:30 p.m.	Speaker 3–17: Bhavya Kailkhura, Lawrence Livermore National Laboratory *Can we design assured deep learning systems?*
12:30–1:00 p.m.	Break

Session Chair: Iris Bahar, Brown University
Session Cochair: Laura Pullum, ORNL

1:00–1:30 p.m.	Invited Speaker 3–8: Auralee Edelen, Stanford/SLAC *Digital twins for particle accelerators at SLAC*
1:30–1:45 p.m.	Speaker 3–18: Xueping Li, University of Tennessee (UTK) *Maintenance Advanced Technology Initiative (MATI)*

1:45–2:00 p.m.	Speaker 3–19: Anthony Corso, Stanford Intelligent Systems Lab *Adaptive stress testing for validating safety-critical autonomous systems*
2:00–2:15 p.m.	Speaker 3–20: Aashwin Mishra, SLAC National Laboratory *Reliable uncertainty quantification for deep learning applications in particle accelerators*
2:15–3:30 p.m.	Breakout Session Out-Briefs
3:00–3:30 p.m.	Breakout
3:30–4:30 p.m.	Breakout Sessions Out-Briefs (continued)
4:30–5:00 p.m.	Workshop Wrap-Up and Next Steps

PRESENTATION INFO

**In order of presentation*

Keynote Presentation

Michael Grieves, Florida Institute of Technology

Title: *Intelligent digital twins: The role of AI and ML in the future of digital twins*

Abstract: Dr. Grieves will discuss how AI and ML will enhance the ability of DTs. He will discuss how DTs will evolve to adopt AI and ML in all aspects of the product lifecycle and share his prediction of the development trajectory that DTs are on.

Bio: Dr. Michael Grieves is an internationally renowned expert in product lifecycle management (PLM) and originated the concept of the DT. His focus is on virtual product development; engineering; systems engineering; complex systems; manufacturing, especially additive manufacturing (AM); and operational sustainment. Dr. Grieves wrote the seminal books on PLM, *Product Lifecycle Management* and *Virtually Perfect: Driving Innovative and Lean Products through PLM*. He has consulted and/or done research at some of the top global organizations, including NASA, Boeing, Newport News Shipbuilding, and General Motors.

Dr. Michael Grieves is currently at the Florida Institute of Technology in Melbourne, Florida, where he helped form the Center for Advanced Manufacturing and Innovative Design. He is currently chief scientist of advanced manufacturing, executive vice president of operations, and interim chief financial officer at the Florida Institute of Technology.

Track 1: Construction of DTs

Nathan Kutz, University of Washington

Title: Targeted use of deep learning for physics and engineering

Abstract: ML and AI algorithms are now being used to automate the discovery of governing physical equations and coordinate systems from measurement data alone. However, positing a universal physical law from data is challenging because (1) an appropriate coordinate system must also be advocated and (2) simultaneously proposing an accompanying discrepancy model to account for the inevitable mismatch between theory and measurements must be considered. A combination of deep learning (DL) and sparse regression, specifically the sparse identification of nonlinear dynamics (SINDy) algorithm, shows how a robust mathematical infrastructure can be formulated for simultaneously learning physics models and their coordinate systems. This can be done with limited data and sensors. The methods are demonstrated on a diverse number of examples by showing how data can maximally be exploited for scientific and engineering applications.

Bio: Nathan Kutz is the Yasuko Endo and Robert Bolles Professor of Applied Mathematics at the University of Washington, having served as chair of the department from 2007 to 2015. He earned his BS in physics and mathematics from the University of Washington in 1990 and a PhD in applied mathematics from Northwestern University in 1994. He was a postdoc in the applied and computational mathematics program at Princeton University before taking his faculty position. He has a wide range of interests, including neuroscience and fluid dynamics, in which he integrates ML with dynamical systems and control.

Farinaz Koushanfar, University of California San Diego

Title: Robust and private machine learning

Abstract: The fourth industrial revolution shaped by ML algorithms is underway. However, the wide-scale adoption of the emerging intelligent learning methodologies is hindered by security, privacy, and safety considerations in sensitive scenarios such as smart transportation, healthcare, warfare, and financial systems. This talk discusses recent progress in devising automated end-to-end algorithms, hardware, and software codesign, optimization, and acceleration of assured ML and privacy preserving systems. A summary of challenges and opportunities ahead is also included.

Bio: Farinaz Koushanfar is the Henry Booker Scholar Professor of Electrical and Computer Engineering at the University of California San Diego, and the founding codirector of the Center for Machine Intelligence, Computing, and Security. She is a well-known leader in automated holistic crosslayer codesign and optimization of ML, security, and privacy-preserving computing. Dr. Koushanfar is a fellow of the IEEE and a fellow of the Kavli Foundation Frontiers of the National Academy of Engineering. She has received a number of awards including the Presidential Early Career Award for Scientists and Engineers from President Obama, the ACM SIGDA Outstanding New Faculty Award, Cisco IoT Security Grand Challenge Award, Massachusetts Institute of Technology (MIT) Technology Review TR-35, Qualcomm Innovation Awards, as well as Young Faculty/CAREER awards from NSF, DARPA, the Office of Naval Research, and the Army Research Office.

Eric Darve, Stanford University

Title: Machine learning for inverse modeling in mechanics

Abstract: The Automatic Differentiation Library for Computational and Mathematical Engineering (ADCME) is a novel computational framework used to solve inverse problems involving physical simulations and deep neural networks (DNNs). By describing physical laws with partial differential equations (PDEs) and substituting unknown components with DNNs, the physics are preserved to the largest extent possible while leveraging DNNs for data-driven modeling. To train the DNNs within a physical system, ADCME expresses both numerical simulations (e.g., finite element methods) and DNNs as computational graphs and calculates the gradients using reverse-mode automatic differentiation. A system of reusable and flexible numerical simulation operators was built to support gradient-backpropagation for many engineering applications, such as seismic inversion, constitutive modeling, and Navier-Stokes equations. ADCME also provides a computational model for conducting large-scale

inverse modeling using MPI (Message Passing Interface) and has been deployed across thousands of cores. The ADCME software is open-source and available at https://github.com/kailaix/ADCME.jl.

Bio: Professor Darve earned his PhD in applied mathematics at the Jacques-Louis Lions Laboratory in the Pierre et Marie Curie University, Paris, France. His advisor was Prof. Olivier Pironneau, and his PhD thesis was entitled "Fast Multipole Methods for Integral Equations in Acoustics and Electromagnetics." He was previously a student at the Ecole Normale Supérieure, rue d'Ulm, Paris, in mathematics and computer science. Prof. Darve became a postdoctoral scholar with Profs. Moin and Pohorille at Stanford University and NASA Ames Research Center in 1999 and joined the faculty at Stanford University in 2001. He is a member of the Institute for Computational and Mathematical Engineering. His research interests include numerical linear algebra, ML for engineering, high-performance, and GPU computing.

Luke Scime, ORNL

Title: Creating scalable digital twins for advanced manufacturing

Abstract: AM promises to revolutionize the manufacturing paradigm by enabling rapid design iteration, fabrication of lightweight components, site-specific microstructure control, and simplified component assembly. However, these 3D printing processes are relatively new, and the community does not yet have the extensive experience needed to understand and control the variation observed in these processes. The use of DTs and augmented intelligence techniques will allow rapid understanding and leveraging of these manufacturing processes now—without waiting decades.

Bio: Luke Scime is an associate staff scientist in the Energy Systems Analytics group at ORNL's Manufacturing Demonstration Facility. Luke earned his PhD in mechanical engineering from Carnegie Mellon University in 2018, and his research focuses on leveraging AI and computer vision techniques for AM.

Rose Yu, University of California San Diego

Title: Physics-guided AI for learning spatiotemporal dynamics

Abstract: Although DL has shown tremendous success in many domains, it remains a grand challenge to incorporate physical principles to such models for applications in physical sciences. This presentation discusses (1) Turbulent-Flow Net—a hybrid approach for predicting turbulent flow by marrying well-established computational fluid dynamics (CFD) techniques with DL—and (2) Equivariant Net—a systematic approach to improve generalization of spatiotemporal models by incorporating symmetries into DNNs. The advantages of these approaches are demonstrated for a variety of physical systems including fluid and traffic dynamics.

Bio: Dr. Rose Yu is an assistant professor at the University of California San Diego. Her research focuses on advancing ML techniques for large-scale spatiotemporal data analysis with applications to sustainability, health, and physical sciences. A particular emphasis of her research is on physics-guided AI, which aims to integrate first-principles with data-driven models.

WaiChing Sun, Columbia University

Title: Microstructure-sensitivity plasticity inferred via graph neural network

Abstract: This talk will provide an overview of a geometric learning framework that builds interpretable macroscopic surrogate elastoplasticity models inferred from subscale direction numerical simulations (DNS) for polycrystalline materials. A graph convolutional neural network is used to deduce low-dimensional descriptors that encode the evolution of particle topology under path-dependent deformation and replace internal variables. To circumvent the lack of interpretability of the classical black-box neural network, a higher-order supervised ML technique is introduced to generate components of elastoplastic models, such as elasticity function, yield function, hardening mechanisms, and plastic flow. The geometrical interpretation in the principal stress space allows the use of convexity and smoothness to ensure thermodynamic consistency. Speed function from the Hamilton-Jacobi equation is deduced from the DNS data to formulate hardening and non-associative plastic flow rules governed by the evolution of the low-dimensional descriptors.

Bio: Dr. WaiChing Sun is an associate professor of civil engineering from Columbia University and a former research scientist at Sandia. His research focuses on data-driven mechanics and plasticity for crystalline, granular, and porous solids across length scales.

George Em Karniadakis, Brown University

Title: DeepM&Mnet: A new neural network architecture based on operator regression for digital twins

Abstract: New NNs are introduced to learn functionals and nonlinear operators from functions and corresponding responses for system identification. The universal approximation theorem of operators suggests the potential of NNs in learning from scattered data in any continuous operator or complex system. First, the theorem is generalized to DNNs and subsequently applied to design (1) a new composite neural network with small generalization error, (2) the deep operator network (DeepONet) consisting of a neural network for encoding the discrete input function space (branch net), and (3) another neural network for encoding the domain of the output functions (trunk net). DeepONet demonstrably learns various explicit operators (e.g., integrals, Laplace transforms, fractional Laplacians) as well as implicit operators that represent deterministic and stochastic differential equations. More generally, DeepONets can learn multiscale operators that span many scales and are trained by diverse data sources simultaneously. Using DeepOnet as a foundation, DeepM&Mnet was designed to use supervised learning with only a few data to simulate complex multiscale and multiphysics systems. DeepM&M is demonstrated for hypersonics problems as well as a multiphysics electroconvection problem.

Bio: George Karniadakis is from Crete. He received his SM and PhD from the Massachusetts Institute of Technology (MIT, 1984/87). He was appointed lecturer in the Department of Mechanical Engineering at MIT and he subsequently joined the Center for Turbulence Research at Stanford University and the NASA Ames Research Center. He joined Princeton University as an assistant professor in the Department of Mechanical and Aerospace Engineering and as associate faculty in the Program of Applied and Computational Mathematics. He was a visiting professor at Caltech in 1993 in the

Aeronautics Department and joined Brown University as an associate professor of applied mathematics in the Center for Fluid Mechanics in 1994. After becoming a full professor in 1996, he has continued as a visiting professor and senior lecturer of ocean/mechanical engineering at MIT. He is an AAAS Fellow (2018–present), Fellow of the Society for Industrial and Applied Mathematics (SIAM, 2010–present), Fellow of the American Physical Society (APS, 2004–present), Fellow of the American Society of Mechanical Engineers (ASME, 2003–present) and Associate Fellow of the American Institute of Aeronautics and Astronautics (AIAA, 2006–present). He received the Alexander von Humboldt award (2017), the Ralf E. Kleinman award from SIAM (2015), the J. Tinsley Oden Medal (2013), and the CFD award (2007) from the US Association in Computational Mechanics. His h-index is 105, and he has been cited over 53,500 times.

Piyush Modi, NVIDIA Corporation

Title: Tools to accelerate design, development, and deployment of digital twins

Abstract: DTs are playing an increasingly important role in modern cyber physical systems by acting as a mirror of the real world to simulate, predict, and optimize operations of industrial assets, systems, and processes. However, building DTs at scale for industrial assets, systems, and processes is a daunting task. This presentation provides examples of tools to facilitate construction of the following:

- Data ingestion, organization, curation, and processing pipeline that helps integrate data that spans the entire lifecycle (i.e., design, manufacturing, operation, and related logistics) of an industrial asset and systems (RAPIDS.ai)
- Ray tracing-aided data visualization platform to render and interact with digital simulations of complex industrial system of systems (e.g., factories with 100s of facilities) (NVIDIA Omniverse)
- Hybrid physics-informed ML models of complex systems (e.g., cumulative damage models of industrial assets) (NVIDIA SimNet, NVIDIA Isaac)
- AI/ML operations to aid collaborative and continuous/federated learning, DT deployment, and management of a secured pipeline at scale (NVIDIA Fleet Command)

These tools will be described in the context of real-world complex cyberphysical systems (e.g., factories, complex industrial asset condition monitoring) to generate a discussion about requirements and their role in accelerating research and development of DTs.

Bio: Piyush Modi is a business development/strategist for the industrial sector at NVIDIA. He is actively engaged with major industrial companies and related research laboratories to conceive and realize industrial AI-enabled solutions. He is interested in real-time DL training/inferencing platforms, architecture, and related algorithms for the industrial use cases spanning inspection, predictive maintenance, and automation.

Prior to NVIDIA, at Sentient.ai as a vice president of engineering, he led a team to build the industry's first distributed AI (DL and Evolutionary Algorithms) platform to serve the model building pipeline needs for e-commerce, financial, and industrial verticals to harness dark compute cycles from 1,000s of GPUs available globally. Over his 20+ year

career, he has held positions of chief technology officer, senior vice president, and head of research laboratory at GE Global Research, BT, Ribbit, IP Unity, and AT&T Bell Labs to deliver major industry-forming outcomes by leveraging speech recognition, VOIP, IoT, and DL technologies. Piyush holds a PhD in electrical and computer engineering (speech recognition) from Rutgers University, an MS in computer science from the University of Tennessee, and a B.Tech. in electrical engineering from the Indian Institute of Technology in Varanasi, India.

David Schmidt, University of Massachusetts Amherst

Title: Accelerated DL representation of turbulent, reacting flow

Abstract: A DT of an engine is only useful if the simulation cost can be reduced to a level at which the representation can fit into a practical design cycle. Typical CFD simulations of turbulent reacting flows provide high-fidelity results at an enormous computational cost. This team is working to accelerate these kinds of multiscale computations through a coarse-graining process that is cognizant of both unresolved physics and numerical errors. The team developed a neural ODE (ordinary differential equation) representation of decaying turbulence and a workflow for coupling DNNs into CFD simulations. For a reacting flow, a two-step process is employed that combines clustering and fitting. Both the turbulence and reaction models fit into a framework that respects basic conservation properties.

Bio: David Schmidt is an SAE Fellow and professor of mechanical engineering from the University of Massachusetts Amherst. His research interests are CFD and ML.

Jeph Wang, LANL

Title: Digital twins for x-ray and neutron cameras

Abstract: X-ray and neutron cameras, as well as cameras for other parts of the electromagnetic spectrum and particles with mass, are widely used in experimental science, medicine, and industry. One of the recent challenges and opportunities is the massive amount of scientific data generated by these devices. Many experiments can easily generate one terabyte of data within a few days. Although camera hardware will advance, as predicted by the Moore's law, image data processing is a significant challenge that may require new automated approaches enabled by AI and ML. Designing DTs for these cameras based on integrating physics principles, materials science, data science, and device engineering and aiming for real-time automated image mining and reduction appears to be an interesting direction that has not yet been thoroughly explored. A broad impact is expected on many areas of fundamental and applied sciences as well as the well-being of society and the environment.

Bio: Jeph Wang is a senior LANL scientist and leads a multidiscipline, multi-institutional collaboration on imaging instrumentation and applications. Recent work includes the billion-pixel x-ray camera (BiPC-X).

Track 2: Application and Deployment of Digital Twins

Felipe Viana, UCF

Title: Digital twins for prognosis applications with hybrid physics-informed neural networks

Abstract: Dr. Viana will challenge the myth that building DTs with ML requires large data sets. First, he will address how physics-driven and data-driven kernels can be combined within DNNs. This framework, which was pioneered in the Probabilistic Mechanics Laboratory at the UCF, allows for a neural network to directly implement differential equations while accounting for uncertainty in the model form as well as observations. Dr. Viana will give an overview of the theoretical aspects and show engineering applications in DTs for failure prognosis of wind turbine main bearings, aircraft fuselage panels, and batteries used to power electric vehicles.

Bio: Dr. Felipe Viana is an assistant professor at UCF, where he leads the Probabilistic Mechanics Laboratory. His research focuses on fusing ML and probabilistic methods with physics-based models for optimization and UQ. Before joining UCF, Dr. Viana was a senior scientist at GE Renewable Energy, where he led the development of computational methods for improving wind turbine performance and reliability. Prior to that role at GE, he spent 5 years at GE Global Research, where he led and conducted research on design and optimization under uncertainty, probabilistic analysis of engineering systems, and services engineering. Dr. Viana holds a PhD in aerospace engineering from the University of Florida and a PhD and MSc in mechanical engineering from the Federal University of Uberlandia, Brazil.

David Stracuzzi, Sandia

Title: Preliminary work on a digital twin for cancer patients

Abstract: The National Cancer Institute and the US Department of Energy (DOE) have recently started a collaboration to develop a DT for cancer patients. Much like DTs for engineered systems, the goal is to model an individual patient (system) in sufficient detail to predict progression of the disease (mechanical flaw), the likely impact of an intervention, or the risk and likely timing of a subsequent adverse event. This talk summarizes one proposed approach to creating a DT for cancer patients, including discussion of several anticipated challenges associated with the data and model construction process. Although the team does not have results on the proposed approach, the described technical issues are likely to be relevant to the development of DTs for engineered systems.

Bio: David Stracuzzi has been studying ML methods for over 20 years, including the past 11 years at Sandia. His recent research emphasizes incorporating domain expertise into models and estimating their predictive uncertainty.

Chetan Kulkarni, KBR, Inc. and NASA Ames Research Center

Title: Hybrid model-based approaches for systems health management and prognostics

Abstract: To facilitate and solve the prediction problem, awareness of the current state and health of the system is key because it is necessary to perform condition-based system health predictions. To accurately predict the future state of any system, one must possess knowledge of its current health state and future operational conditions. Recent achievements of data-driven algorithms in regression of complex nonlinear functions and classification tasks have generated a growing interest in AI for industrial applications. Complex multiphysics models as well as DTs—once purely built on physics and corresponding and simplified lumped-parameter iterations—can now benefit from ML algorithms to mitigate the lack of understanding of some complex behavior. Given models of the current and future system behavior, a general approach of model-based prognostics can solve the prediction problem and further decision making. In principle, data-driven approaches can replace expensive experimental test setups and reduce the number of simulations needed for exploration (e.g., the parametric space of a multiparameter model). Nonetheless, the limitations of pure data-driven methods came to light rather quickly, at least for some industries. In many industrial areas, data acquisition is costly, and the volume of data that can be collected does not satisfy the requirements for an effective model training and cross-validation. Therefore, some recent work in ML focuses on blending physics with data-driven algorithms, thus mitigating the drawbacks of the two approaches and emphasizing respective advantages. Partial physical knowledge of the problem can aid the learning process by "guiding" the algorithm toward solutions that always satisfy the physics driving the system behavior. The result is a hybrid modeling approach combining physical knowledge as well data driven approaches. A hybrid framework for fusing information from physics-based performance models along with DL algorithms for prognostics of complex safety-critical systems is presented. In this framework, physics-based performance models infer unobservable model parameters related to the system components' health, thereby solving a calibration problem in the DL approach.

Bio: Chetan S. Kulkarni is a staff researcher at the Prognostics Center of Excellence, Intelligent Systems Division at the NASA Ames Research Center. He leads a team of researchers in systems health monitoring and prognostics for aeronautics and aerospace applications with a primary area of research in future electric UAVs and aircraft.

Sandra Biedron, University of New Mexico and Element Aero

Title: Experiences in dynamic systems—how a better model can help us understand and control intelligently

Abstract: The team's activities center around dynamic systems, predominantly for scientific inquiry, with interest in the physics-informed construction and use of DTs in real-time control systems. Why? Complex systems can have millions of process variables, change over time, and the subsystems can influence each other. Furthermore, on top of controlling these systems and understanding anomalies/prognostics (e.g., a component failing), the team also wants to analyze them in near real time. For example, in one immediate project funded by EPSCoR, the team wants to analyze the material

properties of what the tool is probing. The team actively uses the Argonne Leadership Computing Facility resources and is establishing a real-time connection between one of these analytical tool systems for control and analysis. The team will soon deploy an edge computing–based subsystem DT at the Facility for Rare Isotope Beams supported by the DOE's Small Business Innovation Research program in Nuclear Physics. Scaling and realization of DL-aided DTs on cloud and HPC systems are of particular interest. Here, a few examples are presented of aspects of these dynamic systems, including an ion-based quantum information science system, particle accelerators, and the precise formation of a 2-CubeSats satellite system. In this way, the team hopes to share the progress thus far and find collaboration opportunities with other community members and workshop attendees to mutually enhance the goals in the deployment of DTs.

The team is also launching a new initiative for dynamic systems led by the University of New Mexico with many collaborators and a major focus on education at all levels. This talk will also touch on much of the foundational research still needed and how that connects to end-use engineering. The proposed institute includes four major thrusts:

(1) Data. How to capture information from temporal, spatial, and streaming data from heterogeneous sources with sparse labeling derived from large-scale, dynamical infrastructures.

(2) Safety. How to provide guarantees and transparency from large-scale learning solutions on high-dimensional dynamic systems.

(3) Explainability. How to capture knowledge from large-scale AI solutions for dynamic systems, thus providing legibility and avenues for human collaboration.

(4) Resource Constrained. How to develop solutions that reduce the power and/or computational requirements, thus enabling learning and adaptability for distributed and dynamic applications.

Bio: Sandra G. Biedron, PhD is a research professor of electrical and computer engineering in the College of Engineering at the University of New Mexico and has served as the chief scientist at Element Aero since 2002. She leads many research projects and recently served as deputy lead engineer for the integration and testing of an innovative naval prototype through a Boeing contract. Formerly, she was the US Department of Defense project office director and a physicist at Argonne National Laboratory and was an associate director of the Argonne Accelerator Institute. Dr. Biedron served as a technical and management consultant on the successful FERMI free-electron laser project at Sincrotrone Trieste (Italy). She is a Fellow of the APS and a Fellow of the SPIE optical society. In 2010, she was presented a Letter of Commendation by the chief of naval research for her technical efforts, and in 2013 she was honored with the George T. Abell Outstanding Mid-Career Faculty Award for the College of Engineering at Colorado State University. In 2018, she received the IEEE Nuclear and Plasma Sciences Society's Particle Accelerator Science and Technology Award. Her interests are many and include particle accelerator systems, laser systems, the use of AI in controls, modelling, and prediction of complex systems, sensors and detectors, and applications of these technologies in science and engineering.

Draguna Vrabie, PNNL

Title: Deep learning digital twins for model predictive control

Abstract: Many real-world systems have unknown dynamics and operate in uncertain environments. Data-driven DL methods offer a pathway to introduce advanced control to complex systems in which physics-based modeling is insufficient. This talk introduces recent work that uses multiple methods to embed domain knowledge in DL representations and trains DL predictive controllers. The talk describes performance comparisons between DL control, traditional model-predictive control, and reinforcement learning (RL) methods on a classical linear time–invariant system. Finally, the talk outlines future research avenues.

Bio: Draguna Vrabie is chief data scientist in the Data Sciences and Machine Intelligence Group, and she serves as team lead for the Autonomous Learning and Reasoning Team at PNNL. Her work at the intersection of control-system theory and ML aims to design adaptive decision and control systems. Her current focus is on DL methodologies and algorithms for design and operation of high-performance cyberphysical systems. Prior to joining PNNL in 2015, she was a senior scientist at United Technologies Research Center in East Hartford, Connecticut. Vrabie holds a doctorate in electrical engineering from the University of Texas at Arlington and an ME and BE in automatic control and computer engineering from Gheorghe Asachi Technical University in Iaşi, Romania.

Junshan Zhang, University of Arizona

Title: Edge intelligence in IoT ecosystems: From continual learning to collaborative learning

Abstract: Many IoT applications demand real-time intelligent decisions. The necessity of real-time edge intelligence dictates that decision making takes place right here, right now at the network edge. Because an edge node often has a limited amount of data and is constrained with computational resources, continual edge learning is advocated to achieve edge intelligence. To this end, the team developed an edge-learning framework in which the edge node learns its model based on local data while leveraging the cloud knowledge transfer or learning from peer edge nodes.

Bio: Junshan Zhang is a professor in the School of Electrical, Computer, and Engineering at Arizona State University. His current research interests are in the general field of information networks and data science.

Hao Huang, GE Research

Title: Industrial data anomaly detection and diagnosis with variable association change

Abstract: Sensors on industrial systems generate multivariate time series, in which each sensor corresponds to one variable. During normal operation, the association (dependency) between variables are mainly stationary. One type of anomaly that is of interest relates to variable association change. Detection and diagnosis of such anomalies refer to pinpointing the time series and the variable associations to which the change relates, which helps in understanding the underlying mechanisms of anomalies. However, detecting such change is difficult because the variable associations are usually unknown

and complicated, and the anomalous samples are usually insufficient for learning the substandard association. This talk presents a neural network that can (1) detect this type of anomaly given multivariate time series as input and (2) locate the association change by learning the nonlinear variable associations from both normal data and the detected anomalies. Specifically, the approach leverages the learned model from normal data to learn the faulty association of the anomalies. Experiments using simulated and real-world industrial data sets show that the model outperforms existing methods.

Bio: Hao Huang is an ML scientist in GE Global Research. He has 10+ years expertise in time series analysis, AD, and diagnosis on industrial data.

Jibonananda Sanyal, ORNL

Title: Transportation/mobility digital twin for Chattanooga

Abstract: The Computational Urban Sciences Group, in partnership with the National Renewable Energy Laboratory and several external stakeholders, have stood up a real-time DT focused on mobility for Chattanooga. The system has brought in 500+ real-time data feeds from 5 systems across 3 institutions, with at least 40 other secondary data sets. This has created an unprecedented opportunity to observe, anticipate, and orchestrate cyberphysical controls toward a 20% energy savings objective for the region. The next phase of this work is expanding the region of interest into Georgia and toward the nearby cities of Nashville and Knoxville, as well as strong collaboration with freight partners and the public works to apply AI-based solutions for transformational changes in transportation efficiency for significant energy savings.

Bio: Jibo leads the activities of the Computational Urban Sciences group at ORNL. His work falls at the intersection of HPC, extreme-scale data and analytics, modeling and simulation, visualization, scalable ML, and sensors and controls for building both research and operational systems focused on complex urban systems at local, regional, and national scales.

Jason St. John, Fermilab

Title: Digital twins for the Fermilab particle accelerator complex

Abstract: A DT was developed to capture the dynamics of the control environment of the gradient magnet power supply (GMPS) for the Fermilab Booster synchrotron. Using the DT as the environment, a multidisciplinary team successfully used RL to train a neural network so it can regulate the GMPS against realistic time-varying perturbations. The final stage of this demonstration will be to deploy the regulator network on a field-programmable gate array (FPGA) and control the accelerator with high precision. This talk outlines the path forward to continuous learning on this FPGA AI GMPS regulator and the open questions faced along the way.

Bio: Jason St. John is a particle physicist who is turning into a data scientist. Officially, he is an applications physicist at Fermilab, where he is involved in projects applying ML to particle accelerators.

Abha Moitra, GE Research

Title: Automating construction of formal assurance case fragments

Abstract: The ever-increasing complexity of cyberphysical systems drives the need for assurance of critical infrastructure and embedded systems. Building assurance cases is a way to increase confidence in systems. In general, the construction of assurance cases is a manual process, and the resulting artifacts are not machine analyzable. The High Assurance Systems team at GE Research is developing technology to support automated generation of formal assurance case fragments for systems, which are both human readable and machine analyzable. These assurance case fragments cover safety and security of systems. The team developed a Semantic Application Design Language Assurance Toolkit (SADL-AT) that includes a semantic model to formalize the goal structuring notation for assurance cases. This presentation describes the SADL-AT and demonstrates the capabilities and effectiveness of SADL-AT by building security and safety assurance case fragments for an unmanned aerial vehicle-based example—a delivery drone. The talk also describes how this approach can be applied to another domain.

Bio: Abha Moitra is a principal scientist at GE Global Research. Her research interests include semantic modeling, knowledge representation, and reasoning as applied to cybersecurity, assurance cases, and manufacturing.

Track 3: Techniques to Provide Assurance

Nurali Virani, GE Research

Title: Humble AI for competence-aware digital twins

Abstract: To safely increase adoption of learned DTs in industry, the team proposes AI approaches that can characterize their own competence and reliability in individual predictions as well as fall back to robust baselines or ask for help when incompetent. The presentation outlines some ideas, results, and challenges in the creation of humble AI and explores how they can be used in industrial and scientific domains.

Bio: Dr. Nurali Virani is a lead scientist in the ML team at GE Research. He is a multidisciplinary researcher and has led several projects, including AI-driven control of wind turbines, AI-driven safe control of power generation gas turbine units, characterizing prediction reliability of ML models, and uncertainty-aware autonomous navigation of ground robots. He was awarded the GE Research CTO Technology Award (5 Under 5) for Outstanding Research in 2018 as well as the 2019 Rudolph Kalman Best Paper Award by ASME. Nurali holds a PhD in mechanical engineering, an MS in electrical engineering, and an MS in mechanical engineering from the Pennsylvania State University. Dr. Virani has 30+ peer-reviewed publications as well as 3 patents.

Jaideep Ray, Sandia

Title: Assembling training data sets for generalizable machine-learned models of physical phenomena

Abstract: Engineering simulations rely on constitutive laws or closure models for small-scale phenomena that are not explicitly modeled. These empirical models are the main source of model-form errors. Of late, these empirical models must be constructed as neural nets and trained on data sets obtained by pooling together a few high-fidelity simulations of moderate complexity. The models have been shown to capture phenomena that eluded engineering simulators in the past.

This talk focuses on how one may assemble the training data set for such ML empirical models so that they may generalize widely. The method relies on (1) being able to cluster the training data set into partitions that each embody a particular type of physics and (2) ensuring diversity among the partitions. The clustering is performed using gaussian mixture models, and the specificity of clusters is limited using information-theoretic criteria. The method is demonstrated on a pool of four DNS (Direct Numerical Simulations) turbulent flow data sets that were used to train a Reynolds-averaged Navier Stokes closure for turbulent stresses.

This method could be used to assemble data sets for training ML models efficiently (i.e., training data sets of tractable sizes that lead to widely generalizable empirical models).

Bio: Jaideep Ray works as a staff engineer at Sandia. His research interests lie in the use of ML and Bayesian calibration in turbulent fluid mechanics and aerothermodynamics (https://www.sandia.gov/~jairay).

Varun Chandola, University at Buffalo

Title: Anomaly detection and clustering for evolving data streams

Abstract: Two salient aspects of complex system behavior are self-organizing and emergent behavior. Monitoring the health of such complex systems requires AD methods that can model the self-organizing or clustering behavior and adapt to evolving and emergent behavior while identifying anomalies. The focus of this presentation is to present an extreme-value, theory-based Bayesian methodology that can identify anomalies and clusters in streaming data without making strict assumptions about the clustering structure and the nature of the anomalies.

Bio: Varun Chandola is an associate professor at the State University of New York at Buffalo (UB) in the Computer Science Department and the Center for Computational and Data-Enabled Science and Engineering. His research covers the application of data mining and ML to problems involving big and complex data and focuses on AD from big and complex data. Before joining UB, he was a scientist in the Computational Sciences and Engineering Division at the ORNL. He has a PhD in computer science and engineering from the University of Minnesota.

Bhavya Kailkhura, Lawrence Livermore National Laboratory

Title: Can we design assured deep learning systems?

Abstract: Mission-critical applications present security and data privacy concerns and demand predictable behavior and strong assurance to achieve safe and correct operation. Unfortunately, modern ML-based systems can be easily hacked because the community lacks the necessary tools to make them foolproof (i.e., obtain guarantees on their robustness and safety). This talk presents recent progress on overcoming this drawback and making DL provably secure. This talk also presents the first open-source PyTorch-compatible library to design foolproof DL models.

Bio: Bhavya Kailkhura is a research staff member at Lawrence Livermore National Laboratory. His research interests are robust ML, signal processing, and optimization. He leads several projects in robust ML systems for high-regret applications.

Auralee Edelen, Stanford, SLAC

Title: Digital twins for particle accelerators at SLAC

Abstract: The controllable settings of particle accelerators often must be adjusted to provide custom charged particle–beam characteristics for different applications or experiments. Simulation models can aid this process, but they are often either too computationally intensive to execute in real time or do not capture the empirical behavior of the accelerator accurately enough for use in control. In addition, myriad sources of uncertainty and changes in accelerator responses over time complicate the modeling process. This presentation provides an overview of progress at the SLAC National Accelerator Laboratory to produce online DTs for its particle accelerators, including deployment on the accelerator control system.

Bio: Auralee is a Panofsky Fellow at the SLAC National Accelerator Laboratory, where she works on developing ML-based approaches for modeling and control of particle accelerators. She arrived at SLAC as a research associate in 2018. During her graduate studies, Auralee worked with Fermilab on early proof-of-principle studies in applying modern neural network–based approaches to particle accelerators. Auralee also has been active in the particle accelerator community for education and promotion of ML, including—for example—helping to organize and provide tutorials for several workshops on ML for particle accelerator applications.

Xueping Li, UTK

Title: Maintenance Advanced Technology Initiative (MATI)

Abstract: MATI is in support of a Plant Directed Research and Development Project at the Y-12 National Security Complex. It focuses on an integrated approach to providing supporting sensor technologies and analytical tools required to support the Consolidated Nuclear Security-wide World Class Maintenance Organization program. The team developed an architecture that uses internet-based sensors and related platforms (e.g., IoT-enabled maintenance management system [IMMS]). IMMS is a network of physical equipment that contains embedded technology to communicate and sense or interact with their states or the surrounding environment in real time. IMMS also includes a suite of

software solutions as a DT embedded with ML algorithms, data visualization, and condition-based maintenance and predictive maintenance models.

Bio: Dr. Xueping Li is a professor in the Department of Industrial and Systems Engineering at UTK. He is the director of the Ideation Laboratory (iLab) and the codirector of the Health Innovation Technology and Simulation Laboratory.

Anthony Corso, Stanford Intelligent Systems Lab

Title: Adaptive stress testing for validating safety-critical autonomous systems

Abstract: Safety-critical autonomous systems require rigorous testing before deployment. Owing to the complexity of modern systems, formal verification may be impossible, and real-world testing may be dangerous and expensive during development. Simulation-based testing is a good alternative but requires the generation of challenging scenarios. Human-designed test cases may not adequately cover the space of possible scenarios and might miss rare or emergent failures. This work presents adaptive stress testing (AST), a technique that uses ML to automatically discover the most likely failures of an autonomous system in simulation. Here, the autonomous system is treated as a black box, and RL is used to manipulate its environment toward challenging or critical scenarios. As demonstrated, the AST can be used for finding failure examples in autonomous driving and aviation domains. Techniques are introduced to improve the scalability of AST to large state spaces and to improve computational efficiency when repeatedly validating related systems.

Bio: Anthony Corso is a 6th-year PhD student in the Aeronautics and Astronautics Department at Stanford University where he is advised by Professor Mykel Kochenderfer in the Stanford Intelligent Systems Laboratory. He studies approaches for the validation of safety-critical autonomous systems with an emphasis on interpretability, scalability, and sample efficiency. His research interests also include RL, optimization, transfer learning, and modeling complex dynamical systems.

Aashwin Mishra, SLAC National Laboratory

Title: Reliable uncertainty quantification for deep learning applications in particle accelerators

Abstract: Particle accelerators find applications in a wide variety of industrial, medical, scientific, and security tasks. Extended time spent on tuning and control of accelerators is expensive because it throttles output. In this context, DNNs are increasingly applied to engender surrogate models and DTs for accelerator applications. However, DL-based models suffer from manifold sources of epistemic and aleatoric uncertainties, are prone to overconfident predictions for out-of-distribution samples and are vulnerable to adversarial attacks. Deployment in high-regret and safety-critical systems, such as particle accelerators, requires reliable measures of predictive uncertainty from DL models. This talk describes the applications of Bayesian Neural Networks (BNNs) to provide accurate predictions with quantified uncertainties for particle accelerator surrogate models. Problems are selected across different designs (e.g., storage rings, beam lines for free electron lasers, the LCLS-II injector) and diverse data volumes and formats. BNN performance is compared to extant approaches in these tasks. Finally, the

presentation describes the calibration of uncertainty estimates from BNNs, such that the interval predictions are reliable.

Bio: Aashwin Mishra is a project scientist in the ML division at SLAC. His work focuses on uncertainty estimation and interpretability for DL models.

APPENDIX C. AIRES 2 WORKSHOP ATTENDEES

FIRST NAME	LAST NAME	ORGANIZATION/AFFILIATION
Ahmedullah	Aziz	University of Tennessee, Knoxville (UTK)
Vittorio	Badalassi	Oak Ridge National Laboratory (ORNL)
Iris	Bahar	Brown University
Matthew	Barone	Sandia National Laboratories (Sandia)
Sylvain	Bernard	Sandia
Sandra	Biedron	University of New Mexico and Element Aero
Willem	Blokland	ORNL
Patrick	Blonigan	Sandia
Thomas	Britton	Jefferson Laboratory
Yu	Cao	Arizona State University
Yanzhao	Cao	Auburn University
Alessandro	Cattaneo	Los Alamos National Laboratory (LANL)
Dave	Caulton	Athena Development
Varun	Chandola	State University of New York at Buffalo
Samrat	Chatterjee	Pacific Northwest National Laboratory (PNNL)
Sanjay	Choudhry	NVIDIA
Eric	Church	US Department of Energy's (DOE's) High-Energy Physics program
Michael	Churchill	Princeton Plasma Physics Laboratory
Jamie	Coble	UTK
Matteo	Corbetta	KBR Inc. and NASA Ames Research Center
Anthony	Corso	Stanford Intelligent Systems Lab at Stanford University
Eric	Darve	Stanford University
Warren	Davis	Sandia
Nathan	DeBardeleben	LANL
Jan	Drgona	PNNL
Eden	Eager	Sandia
Auralee	Edelen	Stanford and SLAC National Accelerator Laboratory
Stephan	Eidenbenz	LANL
John	Emery	Sandia
David	Etim	National Nuclear Security Administration
Katherine	Evans	ORNL
Mariana	Fazio	University of New Mexico
Hal	Finkel	DOE Advanced Scientific Computing Research
Garrison	Flynn	LANL
Michael	Ford	Argonne National Laboratory (Argonne)

FIRST NAME	LAST NAME	ORGANIZATION/AFFILIATION
Ari	Frankel	Sandia
Christopher	Garasi	Sandia
Gerald	Geernaert	DOE
Brian	Giera	Lawrence Livermore National Laboratory (LLNL)
Maria	Glenski	PNNL
Humberto	Godinez	LANL
Michael	Grieves	Florida Institute of Technology
Aric	Hagberg	LANL
Zhizhong	Han	The University of Maryland, College Park
Adi	Hanuka	SLAC National Laboratory and Stanford
Dirk	Hartmann	Siemens
Christy	Hembree	ORNL
Jason	Hick	LANL
Nicki	Hickmon	Argonne
Jacob	Hinkle	ORNL
Forrest	Hoffman	ORNL
Hao	Huang	GE Research
Andy	Huang	Sandia
Amy	Hull	Nuclear Regulatory Commission (NRC)
Raj	Iyengar	NRC
Milan	Jain	PNNL
Prashant	Jain	ORNL
Zheming	Jin	ORNL
Reese	Jones	Sandia
Jessica	Jones	Sandia
Daniel	Ju	NRC
Diana	Kafkes	Fermi National Accelerator Laboratory (Fermilab)
Bhavya	Kailkhura	LLNL
George	Karniadakis	Brown University
Satish	Karra	LANL
Kishore	Kasichainula	Arizona State University
Saad	Khairallah	LLNL
Farinaz	Koushanfar	University of California San Diego
Sharlotte	Kramer	Sandia
Gokul	Krishnan	Arizona State University
Chetan	Kulkarni	KBR Inc. and NASA Ames Research Center
Achalesh	Kumar	GE Research
Thomas	Kurfess	ORNL

FIRST NAME	LAST NAME	ORGANIZATION/AFFILIATION
Nathan	Kutz	University of Washington
Ajay	Kuzhively	Arizona State University
Jonathan	Kyle	Ansys Inc.
Michael	Lang	LANL
Kody	Law	University of Manchester
Bob	Ledoux	Advanced Research Projects Agency–Energy (ARPA-E)
Steven	Lee	DOE Advanced Scientific Computing Research (ASCR)
Angeline	Lee	LLNL
Xueping	Li	UTK
Steve	Lidia	Facility for Rare Isotope Beams and Michigan State University
Frank	Liu	ORNL
Dan	Lu	ORNL
Massimiliano	Lupo Pasini	ORNL
Barney	Maccabe	ORNL
Mahboubeh	Madadi	San Jose State University
Yiorgos	Makris	University of Texas Dallas
Shah	Malik	NRC
Carianne	Martinez	Sandia
David	Mascarenas	LANL – Engineering Institute
Christopher	Mayes	SLAC National Accelerator Laboratory
Craig	Miller	Ansys Inc.
Aashwin	Mishra	SLAC National Accelerator Laboratory
Piyush	Modi	NVIDIA
Arvind	Mohan	LANL
Abha	Moitra	GE Research
Curt	Nehrkorn	ARPA-E/Booz Allen Hamilton
David	Najera	ATA Engineering
Justin	Newcomer	Sandia
Ron	Oldfield	Sandia
Jonathan	Ozik	Argonne
Pinaki	Pal	Argonne
Vincent	Paquit	ORNL
Michael	Parks	Sandia
Jyotishman	Pathak	Cornell University
Laura	Pullum	ORNL
Mihaela	Quirk	National Nuclear Security Administration
Pradeep	Ramuhalli	ORNL

FIRST NAME	LAST NAME	ORGANIZATION/AFFILIATION
Vivek	Rathod	ORNL
Daniel	Ratner	SLAC National Accelerator Laboratory
Satha	Raveendra	ESI Group
Jaideep	Ray	Sandia
Pedro Alberto	Resendiz Lira	LANL
Juan	Restrepo	ORNL
Ahmad	Rushdi	Sandia
Nandakishore	Santhi	LANL
Jibonananda	Sanyal	ORNL
Mina	Sartipi	University of Tennessee at Chattanooga
Abhinav	Saxena	GE Research
David	Schmidt	University of Massachusetts Amherst
Malachi	Schram	PNNL
Luke	Scime	ORNL
Tom	Scripter	Sandia
Phillip	Scruggs	Consolidated Nuclear Security
Sudip	Seal	ORNL
Jenifer	Shafer	DOE and ARPA-E
Steven	Sloman	Brown University
Matthew	Smith	Sandia
Vladimir	Sobes	UTK
Sibendu	Som	Argonne
Salvador	Sosa	University of New Mexico
Brian	Spears	LLNL
Bill	Spotz	DOE
Sarat	Sreepathi	ORNL
Jason	St. John	Fermilab
Grant	Stewart	LANL
Panos	Stinis	PNNL
David	Stracuzzi	Sandia
Jan	Strube	PNNL
WaiChing	Sun	Columbia University
Christine	Sweeney	LANL
Daniel	Tartakovsky	Stanford University
Chris	Tennant	Jefferson Laboratory
Hoang	Tran	ORNL
Nathaniel	Trask	Sandia
John	Turner	ORNL

FIRST NAME	LAST NAME	ORGANIZATION/AFFILIATION
Kristopher	Velazquez	Lockheed Martin
Velimir	Vesselinov	LANL
Felipe	Viana	University of Central Florida
John	Vickers	NASA
Nurali	Virani	GE Research
Svitlana	Volkova	PNNL
Draguna	Vrabie	PNNL
Sarma	Vrudhula	Arizona State University
Adam	Wachtor	LANL – Engineering Institute
Jeph	Wang	LANL
Dali	Wang	ORNL
Haoyu	Wang	Argonne
Hong	Wang	ORNL
Justin	Weinmeister	ORNL
David	Womble	ORNL
Mitchell	Wood	Sandia
Vaibhav	Yadav	Idaho National Laboratory
Srikanth	Yoginath	ORNL
Huafeng	Yu	Boeing
Rose	Yu	University of California San Diego
Guannan	Zhang	ORNL
Jiaxin	Zhang	ORNL
Junshan	Zhang	Arizona State University
Shukui	Zhang	Jefferson Laboratory