Federal Register Notice: 89 FR 51554, [Federal Register :: Networking and Information Technology Research and Development Request for Information on Digital Twins Research and Development], June 18, 2024.

Request for Information on the National Digital Twins R&D Strategic Plan

Sandia National Laboratories

Sandia National Laboratories' Response to the Networking and Information Technology Research and Development Request for Information on Digital Twins Research and Development

Date: July 26, 2024
Organization Name: Sandia National Laboratories, [redacted]

POC: Anna Wakeland; Technical Business Development Specialist; [redacted]

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. [SAND2024 - 09657O]

Executive Summary:

Sandia recommends emphasis on the following R&D topic areas for strategic planning:

- AI, as it pertains to the construction of digital twins, data analysis, fast-running surrogate models, and real-time feedback control of manufacturing processes,
- Data management practices related to data collected from physical systems or physical systems augmented with machine learned models,
- Establish ecosystems founded on interagency collaboration and composed of public-private consortia that develop and promote standards, and fund interdisciplinary work that addresses identified gaps,
- International collaborations like the Digital Twin Consortium and the Omniverse for Microelectronics Fabrication,
- Long-term research investments in the bidirectional flow of information needed to construct digital twins, bringing together experimental and digital computing technologies, and combining models of various fidelity,
- Methods for creating safe and secure digital twins based on rigorous model-based engineering concepts,
- And verification and validation methods stemming from complex systems analysis

***Artificial Intelligence (AI)*:** AI and Digital Twins:

The Office of the Secretary of Defense in 2016 defined a future Digital Engineering Ecosystem for the Department of Defense as "the interconnected infrastructure, environment, and methodology (process, methods, and tools) used to store, access, analyze, and visualize evolving systems/data and models to address the needs of the stakeholder." Similarly, we are constructing a Digital Engineering Ecosystem for Nuclear Deterrence (ND) at Sandia that enables advanced model-centric workflows and applications across engineering domains and the lifecycle of weapon systems. Moving away from static document-based processes towards a dynamic model-based ecosystem will enable the creation of an authoritative digital source of truth for each weapon program. This will ensure the right information is always available for accelerated decision making and accurate and appropriate assessment across the weapon lifecycle. This model is applicable to other high consequence areas such as satellites, global security and other energy related areas.

Digital twins (DTs) are real-time virtual renderings of the physical world. A primary capability offered by a DT is rapid access to - and communication around - data and information. With the high consequence complex systems within the national security enterprise, most teams are challenged to manage large volumes of technical documentation, often in multiple formats and duplicated throughout the organization, over multiple decades leading to fragmentation and degradation of information integrity. This leads to significant financial and safety risks during operations. DTs can help ensure that information is easy to find when it is needed as well as serving as a single source of truth that supports communication and collaboration by presenting information in context or data pedigree. An effective DT is based on the inherent data structure and governance that maintains the validity and accuracy of the twin. AI connects the digital and physical worlds transforming science and engineering. Enabling real-time virtual representations of the manufacturing process provides production teams with the ability to support faster, smarter, and more cost-effective decision making. Accurately captured models in AI can deepen manufacturers' understanding of complex physical systems and production operations, optimize production scheduling, or simulate "what-if" scenarios to understand the impact of new product introductions, which are critical for responsive national security needs.

AI connects the digital and the physical worlds transforming science and engineering. AI understood methods have the potential to revolutionize the formulation and use of DTs. For example, graphical models like dynamic Bayesian networks can be used to augment equations of physical systems to enable robust updates of DTs at scale [Kapteyn, 2021]. Similarly, AI models can be used to entirely replace physical system models in applications where either the physical system is poorly or where such models are not practical (e.g. modeling human behavior). E.g., Sandia has explored the use of process models and hidden Markov models as DTs to model and analyze human behavior in physical [SAND2023-14702, SAND2024-00219] and cyber systems [KTR-2024-003, SAND2023-06726C]. Many research questions remain concerning the use of AI models as DTs. For example:

- How do we develop trustworthy AI models from sparse, noisy data sources?
- How can domain information, including subject matter expertise, be effectively integrated into AI models?
- How can AI models capture both statistical and logical relationships in physical systems?

Machine-learning and AI will play an integral part in the successful construction of DTs. This need arises from two separate issues. DTs will need to be continuously calibrated to measurements streaming from the physical twin. Continuous data assimilation algorithms e.g., Ensemble Kalman Filters exist, but require the DT to execute quickly on a computer – but for complex systems, this is hardly feasible. Fast-running AI proxies of computationally expensive models will be necessary to make DTs a success; indeed, such DTs are already being used to design inertial confinement fusion targets [Wang et al, 2024] and could be used to design processes to allow fusion power generation. The second use arises from needing the DT to be of a fidelity sufficiently high such that it can successfully exploit and assimilate the information content in the continuous measurements of the DT. All DTs will likely contain phenomenological models / closures that are "curve fits" to historical data; they can (and should) be replaced by AI models trained on data from high-fidelity models that already exist (the feasibility of this process has been widely demonstrated [Duraisamy, 2021]). In many cases, a dense observation of the physical twin may be impossible (due to size/power/weight/accessibility restrictions on the sensors), and it may be necessary to "fill in" the gaps in observational data, for visualization, interpretability, fast anomaly detection and classification of the observed phenomena and perhaps also for decision-making and control of the physical twin. "Filling in the blanks" (or conditional generation) is a prototypical generative AI problem, and efforts have already begun to address scientific problems via the construction of generative (or "super-resolution") models [Deng et al, 2019, Fukami et al, 2021]

At Sandia, we have demonstrated that by using a machine learned anomaly detection approach, it is possible to detect voids and other defects in materials DT. This paves the way for future research in integrating materials DT with its physical counterpart. Detecting anomalies in fatigued and fractured experimental materials is an interesting yet challenging topic. The reasons are threefold. First, the anomalous microstructure feature that gives rise to structural failure is small, sometimes in the order of 10^{-7} of the interrogated volume. This, in turn, results in a highly imbalanced classification problem in machine learning (ML). Second, the consequence is high, in the sense that the test specimen is destructed in such case. Third, the convolution between microstructure stochasticity and the small probability of void nucleation, growth, and coalescence makes failure and fracture a hard-to-predict and challenging problem in materials science due to its irreproducibility, even experimentally. In [Tran, 2024] we developed a materials DT and applied anomaly detection methods to detect voids and anomaly in additive manufacturing (AM). The materials DT is driven by two integrated computational materials engineering (ICME) models, which are kinetic Monte Carlo (kMC) and crystal plasticity finite element method (CPFEM).

*Investigate Task Suitability and Sustainment of AI Tools***:** This action will involve locating potential case studies that examine the suitability of AI tools for use in DT application. The objective is to identify and gather information on real-world examples of AI applications, focusing on both successful and unsuccessful implementations. The objective is to distinguish operations currently identified as benefitting from AI augmentation from those that should remain under human control, considering factors such as complexity, ethical implications, decision-making requirements, and the potential for AI to enhance or detract from the task. Exercises will also be undertaken to evaluate hypothetical near-future operations and AI capabilities in the same manner. Through this investigation, a clear framework will be established to guide the desired use levels and roles of AI in supporting human operators. It aims to define the tasks that are appropriate for AI facilitation as determined by human systems researchers and the DoD and to highlight areas where the current limitations of AI technology and practical challenges make its application inadvisable.

Develop Requirements for High-Consequence Human-AI Decision-Making: Extending the activities of #1, this proposed action will evaluate the tasks deemed suitable for AI use from a risk standpoint. These tasks will be assessed across the spectrum of error consequence severity. Our objective will be to more fully characterize how consequence severity should be accounted for in deciding when and how to utilize AI tools in DTs. We will use formal methods such as failure modes and effects analyses (FMEAs) and less formal "what-if" scenario case studies to explore the consequence gradient with respect to potential AI-assisted decision-making activities.

Data: Encourage Adoption of Data Management Best Practices:

Digitals Twins will likely be purely data-driven or hybrid (a combination of approximate physics models augmented with machine-learned corrections). This is because the DT will need to be computationally efficient so that it can be kept calibrated to observations from the physical twin. This, in turn, implies that a DT will be associated with two data corpora – a training dataset (TD) for its machine-learned components and a calibration dataset (CD) gathered from the physical twin. Their (potentially real-time) integration with a DT demands new interoperability standards that enable seamless data exchange across models, systems, and platforms.

Strategic challenges: There are enormous challenges in assembling the TD, as the data must be informative about the physical processes that are of concern to the physical twin (e.g., aging, failures, tampering/sabotage etc.) *and* can be observed via sensors. While one can generate the TD using high-fidelity models, there are no automated means to detect whether a proposed dataset has the correct physics, in the operational regimes of interest. There have been attempts to gauge the physics-content of TDs, but these efforts are in their infancy [Barone et al, 2022]; without a well-developed capability to do so, assembling TDs for realistic systems, especially for rare phenomena like failures (i.e., the data will be largely model-generated), seem to be infeasible. A second, but related challenge lies in making such datasets FAIR (Findable, Accessible, Interoperable and Reusable). Finding datasets implies being able to index them in some fashion, and to date, only textual (and perhaps pictorial) descriptions of a dataset can be indexed. For quantitative modeling purposes, these are insufficient; instead, one needs summaries of the dataset that not only provide the physics content of latent information in the TD, but also (summaries of) features of the data – spatiotemporal autocorrelations, coherent structures, etc., that can be checked for correctness and interpreted via the laws of physics. Prototypes [Bien et al, 2011], which are elements of the TD selected from physically interpretable clustering of the TD [Barone et al, 2022] are one way of summarizing a dataset in this fashion; reducing the dataset into a simplified, machine-learned dynamical system e.g., Universal Differential Equations [Rackauckas et al, 2020] is an another. These methods, rooted in machine-learning, are far from being mature, but are necessary for FAIR datasets. Integrating these representational techniques for TDs into a searchable storage system remains a distant vision.

Tactical challenges: Before TDs can be assembled, they must be generated and contributed to a repository for checking, indexing, and archiving. This implies the need for massive *fast* storage in the computational centers where they are generated and where the DTs are learned. Cloud storage is sufficiently large but too slow to be coupled to supercomputing centers. In addition, a dataset, once accepted into a TD, needs to be described in text and "logged" into a metadata directory for easy perusal. There are efforts to design an appropriate scheme [Gebru et al, 2021] but it is debatable whether these schemata are complete for scientific TD. Further, there does not seem to be any effort to develop such summary textual descriptions for TD for DTs, let alone archive them and make them searchable. Such a capability would not be difficult to construct (i.e., would not require much research), but would need concerted (implementation) effort.

Ecosystem: Establish a National Digital Twin R&D Ecosystem:

To establish a National Digital Twin R&D Ecosystem, we recommend the following:

- Encourage interagency (federal agencies, national labs, industry and academia) collaboration & task force to coordinate DT R&D efforts and organize regular interagency workshops and conferences to share progress, challenges, and best practices.
- Identify the critical research gaps & fund interdisciplinary research projects that address these gaps, focusing on areas such as data integration, real-time analytics, and predictive modeling.

- Develop public-private consortia to collaborate on high-impact DT projects where industry and academia can collaborate on developing and testing DT solutions.
- Develop and promote standards to ensure interoperability and scalability of DT systems.
- Biomedical sciences: Develop DTs of (i) virtual patient model for better decision making / decision support system; (ii) monitor chronic disease to predict future conditions in advance and treat patient on time; (iii) surgical instruments and procedure to enhance preoperative planning and train surgeons, bio-surveillance to monitor public health; (iv) biological systems to simulate drug interactions, efficacy of drug, and identify side effects.
- Common mathematical, statistical, and computational foundations: (i) develop a framework that support multiscale modeling (enabling the integration of models at different spatial and temporal scales); (ii) combine deterministic and stochastic modeling techniques to capture both predictable and random behaviors in complex systems; (iii) use Bayesian inference methods to update DT models in real-time based new data; (iv) deep learning techniques to create highly accurate and scalable DT models; (v) integrate cloud computing technologies with HPC resources to provide flexible, on-demand access to computational power.

International: International Collaborations on Digital Twins:

By bringing together expertise from different countries and sectors, these initiatives are driving innovation and setting the stage for the widespread adoption of DT technology across various domains.

Digital Twin Consortium (https://www.digitaltwinconsortium.org/)

An international consortium including companies including Microsoft, Dell, and GE Digital, along with academic and governmental organizations. The Digital Twin Consortium aims to accelerate the adoption of DT technology across industries by developing standards, guidelines, and best practices. The consortium includes members from around the world, fostering global collaboration and knowledge sharing.

Horizon Europe Digital Twins (https://research-and-innovation.ec.europa.eu/funding/funding-opportunities/funding-programmes-and-open-calls/horizon-europe_en)

The European Union’s Horizon Europe Program is investing in multiple DT activities. These include ocean and earth models as well as models focused on infrastructure and manufacturing.

The International Data Spaces Association (IDSA) (https://internationaldataspaces.org)

IDSA includes members from various countries, including Germany, the USA, Japan, and others. IDSA develops standards and architectures for secure data exchange in DT applications, particularly in manufacturing and logistics. The association’s global membership fosters international collaboration on data standards and interoperability.

Omniverse for Microelectronics Fabrication (https://resources.nvidia.com/en-us-industrial-sector-resources-mc/en-us-industrial-sector-resources/gtc24-s62610?ncid=no-ncid)

This is an industrial collaboration Nvidia, Siemens, TSMC, Samsung and others. Nvidia’s Omniverse capability for digital twinning is being specifically developed to support twins of microelectronics fabs. There is a potential for this to become a standardized tool within this space although there will be competing technologies. This is a capability that could be highly aligned with the CHIPS and Science Act as well as facilitating the incorporation of AI technologies through Nvidia’s other products.

Long Term: Identify Long Term Research Investments:

The bidirectional flows of information required to build functional DTs require a commitment by all parties involved to incorporate a high degree of digital cooperation. The continuous integration and continuous deployment (CI/CD) technologies of software development (e.g. GitHub/GitLab/runners, etc.) are good examples of how such processes are currently developed and maintained. However, DTs require communication between many distinct software and data repositories with proper information protections in place that are both sufficiently secure but yet open to all appropriate entities and containing access to sufficient meta-data that automated processes can be built and maintained. A significant challenge to such a broad based and usable information environment is to provide overarching and broad-based data availability requirements so that the needed data and associated personnel expertise can realistically be supported. Combining experimental software and digital simulations requires a great degree of commitment to cross-disciplinary cooperation.

From a research point of view the main challenges to the development of DTs are the inherent roadblocks arising from scientific and engineering funding structures and culture that are misaligned with the DT ecological viewpoint. Research should be funded to demonstrate specific instances of DT systems that bring together experimental and digital computing technologies and be required to show measurable benefits or failures so as to expose the essential issues needing improvement. Proposals that provide for delivery of multi-domain cross-cutting glue code that could be open sourced should be carefully considered. At this early stage working demonstrations of real systems that expose the essential requirements of DTs should be funded and the results presented as widely as possible. The development of systems and processes that enable easier maintenance and evolution are essential since the whole idea of a DT is a continuous monitoring and improvement.

Rudimentary DTs (per the NITRD definition), with bi-directional coupling, already exist – a common example would be a model of an oil/gas reservoir that is kept calibrated to monitoring wells' data and periodic seismic surveys and which are used to make decisions e.g., enhanced oil recovery, that significantly impact the nature and response of the physical twin. The primary challenge has been the steady improvement in the density, quality and modalities by which measurements of the physical twin are made, as they require a concomitant improvement in the sophistication of the model – one cannot assimilate measurements of a physical process unless it is included in the DT. It is never very clear which processes should be included in the model, and their inclusion invariably increases the computational cost and numerical pathologies of the model (including the embedding of machine-learned phenomenological models / closures for higher fidelity); despite these complexities, the DT must be subjected to sequential data assimilation / calibration to streaming observations. The continuously assimilation of multiscale measurements has been achieved, to some degree, in Earth system models, but employing a hierarchy of model of variable fidelities, but what this hierarchy may be for an arbitrary physical twin is unknown and has not been attempted. Successful sequential data assimilation will require robust calibration techniques that adapt to increasing complex models, or their hierarchies, and these simply do not exist today.

Trustworthy: Realize Secure and Trustworthy Digital Twins:

A key opportunity for DTs is that they can enable comprehensive analyses of security and cyber resilience that are not practical on the physical counterpart. Sandia has previously explored the use of process models as DTs to model and analyze activities in physical [SAND2023-14702, SAND2024-00219] and cyber systems [KTR-2024-003]. A process model is a graphical or logical representation of an operational process or workflow, including events or activities that occur in the workflow, how they are executed, and their logical relationship to other activities and resources. Although there is a robust literature on the use of process models in both the business and engineering applications, many research questions remain concerning their use as DTs. For example:

- How do we develop trustworthy process models from sparse, noisy data sources?
- How can domain information, including subject matter expertise, be effectively integrated into process models?
- How can process models capture both statistical and logical relationships in physical systems?

Ongoing research at Sandia is starting to consider research questions like these (e.g. [SAND2023-06726C]).

DTs seek to represent the behavior of complex cyber-physical systems, thereby inheriting the complexity of the systems themselves. Due to their nonlinearity and vast state spaces, establishing trust in DTs is generically very difficult, especially if they are constructed as monolithic, fully detailed models. R&D is needed to discover appropriate multi-fidelity modeling abstractions and decompositions that make DTs tractable and scalable for analysis and verification.

Trustworthy DTs should ideally be the result of rigorous model-based engineering. That is, models should be constructed iteratively throughout the engineering process and be used to guide *design for analyzability*. This will improve robustness of both the DT and the system itself. It is much easier to establish trust when starting from simpler high-level models and incrementally refining them, in a way that generates mathematical evidence (such as formal verification) of their consistency and conformance.

When the iterative engineering process is complete, the collection of these interrelated models constitutes a more useful DT for the system because it captures the decisions and reasoning that tie the system's components together, and tie the fully detailed behavior to the higher-level requirements and functional models. This allows particular questions about the system to be addressed using DT models of an appropriate scope and detail for the task, providing an improved basis for trust.

VVUQ: Develop Rigorous Methods for Verification, Validation, and Uncertainty Quantification for Digital Twins:

VVUQ represents a foundational commitment to quality and continuous improvement. This means that the DT ecosystems will be continuously changing, improving and evolving. A healthy VVUQ ethos should be expected and funded from the beginning, and embedded into all stages of the DT life cycle, e.g., following the ASME VVUQ Standards.

Complex systems are ubiquitous across all of science and engineering and beyond. Historically, improving understanding of complex systems has involved building models of these systems and using these models to develop a simulation framework; that is, a virtual representation of the system. These simulation frameworks have evolved over the past decade into the concept of a DT [National Academies, 2023], which involves bridging the virtual system and the physical system it represents. The recent SIAM Report on the Future of Computational Science [Hendrickson, 2024] notes that while DTs present enormous opportunities, their development faces great mathematical, statistical, and computational challenges. Models of complex systems are subject to a wide variety of sources of both epistemic and aleatory uncertainty that must be quantified to enable robust and informed decision making. The accurate prediction of the behavior of complex systems, utilizing DTs, are necessary to inform critical decisions that can enhance national security and/or avoid substantial human and financial losses. Moreover, risk assessment is needed to quantify the effect of uncertainties on the severity of predicted outcomes. Rational decision making requires robust, accurate, and computationally efficient methods that can compute quantitative metrics on actionable time scales.

However, as noted in [National Academies, 2023], "verification, validation, and uncertainty quantification as essential tasks for the responsible development, implementation, monitoring, and sustainability of digital twins", however, "a gap exists between the class of problems that has been considered in traditional modeling and simulation settings and the UQ problems that will arise for digital twins." Thus, the integration of robust metrics in decision making workflows faces several critical challenges. First, probabilities must be conditioned on available observational data. While numerous advances have been made in Bayesian inference to characterize epistemic uncertainty, far less attention has been paid to aleatoric uncertainty. Moreover, almost no attention has been given to conditioning both epistemic and aleatoric uncertainty on data in a unified framework. Second, quantifying metrics from the push-forward of posterior distributions is computationally demanding for existing methods that rely solely on high-fidelity simulations. Techniques such as Markov Chain Monte Carlo (MCMC) require numerous evaluations of the simulation model, which are often intractable even on leadership-class computing resources. Third, there is a need for more efficient methodologies for optimal experimental design (OED) to guide data acquisition efforts that will optimally improve model parameter characterization and the subsequent reliability of the predictions made using the computational model. All of these challenges are amplified when models are parameterized by large numbers of uncertain variables and risk-assessment must be executed on actionable time-scales.

Addressing these challenges will require advances on multiple fronts and combining concepts from different fields. Enabling fast and credible predictions for decision making using models of complex systems requires the integration of state-of-the-art UQ methods and Scientific Machine Learning (SciML) approaches. SciML is well positioned to address the high-dimensionality of parameterizations and the computational cost of coupled models that currently limits the complete adoption of UQ in scientific workflows. To serve as the basis for decision-making, new AI/ML approaches must be developed that respect the data, respect the physics and their well-developed numerical treatments, and be able to quantify the uncertainty in their outcomes.

For DTs of complex systems to be truly useful, a holistic UQ framework that addresses both epistemic and aleatoric sources of uncertainty is required. Moreover, such a framework should be replete with computational diagnostics that allow for AI-enabled inferences and decision making in the presence of such uncertainties. Recent work indicates that leveraging information from a population of assets through the solution of an aleatoric stochastic inverse problem to build population-informed priors can significantly reduce the uncertainty in asset-specific Bayesian (epistemic) inferences [White, 2024].

AI system security capability (EXCALIBUR) provides methods to examine (attack and measure) the security of systems that use data-driven models. EXCALIBUR examines broader perspectives beyond the data-driven model and data focusing on (1) the reliability and robustness of the system and how it will behave in novel conditions, (2) the vulnerability of the system to adversarial attacks, and (3) the vulnerability of the data driven model to being affected through cyber-means. EXCALIBUR bridges the gap between established capabilities in machine learning, adversarial machine learning, and cyber-security to focus on AI systems. We provide real-world insights and lessons learned through assessments on actual AI systems including a DT cyber defense technology for process modeling (PROM) that utilized subspace identification techniques to model dynamic behavior in industrial processes. With this "digital twin" of the physical industrial process plant, PROM can operate as a cyberphysical intrusion detection system to detect anomalous behavior. EXCALIBUR evaluated adversarial threats to this DT system and identified weaknesses that may be exploited to compromise system performance and trustworthiness.

Workforce: Cultivate Workforce and Training to Advance Digital Twin R&D:

Integration of AI and DTs in high consequence applications required a specialized workforce with a diverse set of skills. The skills can be broadly categorized into technical expertise, domain expertise, multi-disciplinary skills, and maturity involving awareness including soft skills such as communication. With regard to technical expertise AI and machine learning skills are required for developing the algorithms that drive the predictive and analytical capabilities of DTs. Individuals proficient in data analysis, model training and algorithm optimization will play a crucial role. Further staff will be required to efficiently handle large volumes of data processing, visualizing and managing simultaneously. Fundamental understand of the data structures and governance would be key to enable the workforce. Domain expertise would be addressed by industry specific experts with deep insights and knowledge of the specific high consequence fields such as nuclear deterrence (design and manufacturing) or global security or satellite systems. Individuals who understand the operational aspects of the complex systems being modeled; play a foundational role as they bring the physical connection to help validate the DT models. System engineers and project managers with multi-disciplinary skills who can take a holistic view of the entire system, ensuring that all aspects (hardware, software, data, and processes) work together effectively. Systems engineers play a crucial role in managing and maintaining a complex interconnected physical system which would apply to the DT model as well. Continuous learning and training for the workforce will enable to team to stay updated with the latest advancements which would be both on the job training as well as formal education. Staff supporting environments where workforce can experiment and learn about DT and AI systems without impacting critical path tasks adversely. Workforce knowledgeable in topics such as safety, security, and other regulations and guidelines as applied to DT and AI systems will be integrated into teams.

Sandia's FORGE ND program aims to revolutionize the onboarding and training process for new employees at Sandia National Labs by integrating artificial intelligence (AI), large language models (LLMs), and the digital environment into an apprenticeship-based learning model. This initiative addresses the lag time between theoretical knowledge and practical application, ensuring that new hires become productive more quickly and efficiently.

References

- Barone, M., J. Ray and S. Domino, "Feature Selection, Clustering, and Prototype Placement for Turbulence Data Sets," *AIAA Journal*, 60(3):1332-1246, 2022.
- Bien, J., and Tibshirani, R., "Prototype selection for interpretable classification," *Annals of Applied Statistics*, Vol.5, No. 4, 2011, pp. 2403–2424. https://doi.org/10.1214/11-AOAS495

- Deng, Z., Chuangxin He, Yingzheng Liu, Kyung Chun Kim; Super-resolution reconstruction of turbulent velocity fields using a generative adversarial network-based artificial intelligence framework. *Physics of Fluids* 1 December 2019; 31 (12): 125111.
- Duraisamy, K. (2021). Perspectives on machine learning-augmented Reynolds-averaged and large eddy simulation models of turbulence. *Physical Review Fluids*, *6*(5), 050504.
- Fukami, K., Fukagata, K., & Taira, K. (2021). Machine-learning-based spatio-temporal super resolution reconstruction of turbulent flows. *Journal of Fluid Mechanics*, *909*, A9.
- Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Iii, H. D., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, *64*(12), 86-92.
- Hendrickson, B. A. Aceves, E. Alhajjar, M. Ba nuelos, D. Brown, K. Devine, Q. Du, O. Ghattas, R. Giles, M. Hal, T. Islam, K. Jordan, L. Lin, A. Pothen, P. Raghavan, R. Schreiber, C. Thalhauser, and A. Wilson. Siam task force report: The future of computational science. SIAM Report, 2024
- Kapteyn, M.G., Pretorius, J.V.R. & Willcox, K.E. A probabilistic graphical model foundation for enabling predictive digital twins at scale. Nat Comput Sci 1, 337–347 (2021).
- KTR-2024-003: Eydenberg, M.S., W. Hart, A. Outkin, C. Diaz, R. Valme and A. Beauchaine (2024). LOMA PARDA: Applications of constrained optimization to infer advanced persistent threat activity from noisy data streams. Tech. rep. Sandia National Laboratories.
- National Academies of Sciences Engineering and Medicine. Foundational research gaps and future directions for digital twins, 2023.
- Rackauckas, C., Yingbo Ma, Julius Martensen, Collin Warner, Kirill Zubov, Rohit Supekar, Dominic Skinner, Ali Ramadhan, and Alan Edelman. Universal differential equations for scientific machine learning. *arXiv preprint arXiv:2001.04385*, 2020.
- SAND2024-00219: Eydenberg, M. S., D. Z. Anderson, C. P. Diaz, W. E. Hart, V. J. Leung, and C. D. Ulmer (2024). Using Process Matching to Detect Patterns in U1a Data. Tech. rep. Sandia National Laboratories.
- SAND2023-06726C: Hart, W. E. and A. Outkin (2023). Constrained inference with Hidden Markov Models. Tech. rep., Sandia National Laboratories.
- SAND2023-14702: Hart, W. E. and V. J. Leung (2023). Process Matching: Methods and Optimization Formulations. Tech. rep. Sandia National Laboratories.
- Tran, A., Max Carlson, Philip Eisenlohr, Hemanth Kolla, and Warren Davis, Anomaly Detection in Materials Digital Twins with Multiscale ICME for Additive Manufacturing. *Integrating Materials and Manufacturing Innovation,* June 2024, https://doi.org/10.1007/s40192-024-00360-8.
- Wang, J., N. Chiang, A. Gillette, J. L. Peterson; A multifidelity Bayesian optimization method for inertial confinement fusion design. *Phys. Plasmas* 1 March 2024; 31 (3): 032706.
- White, R., J. Jakeman, T. Wildey, T. Butler. Building Population-Informed Priors for Bayesian Inference Using Data-Consistent Stochastic Inversion. *zrXiv preprint arXiv:2407.13814,* 2024