



DECIDE

Workshop Report
June 8, 2025

DECIDE

Enhancing National Security and Resilience through
Advanced Computational Decision Support



Executive Summary

The DECIDE workshop, held on March 4-5, 2025, brought together leading researchers, U.S. government officials, and industry experts to address the critical need for computational decision-making methods for high-consequence, complex problems. The escalating complexity and interconnectedness of modern infrastructure and ecosystems require sophisticated computational decision support capabilities for timely and informed decision-making in emergencies and high-stakes scenarios. Improved decision support capabilities would ultimately save lives, protect property and critical infrastructure, safeguard public health, and enhance national security and economic strength. The DECIDE workshop was motivated by the vision of the National Strategic Computing Reserve (NSCR), a strategic national resource to provide compute/data resources and bring expertise to bear in times of crisis to help save lives, property, public health, and safety, or to lessen or avert the threat of a catastrophe. Specifically, the workshop recognized the need for an NSCR, advanced the formation of an interdisciplinary stakeholder community, and established a shared view of the opportunities and gaps for a national capability supporting local front-line response to emergencies.

The Imperative for Advanced Computational Decision Support

Our nation confronts a growing number of high-consequence events, including cyberattacks, pandemics, extreme weather, and infrastructure failures. In this era of deeply interconnected and complex systems, timely insight into threats and the coordination of decisions can greatly improve results across jurisdictions, as the impacts of events rapidly cascade across regions. The United States possesses unparalleled computational and data capabilities across its federal agencies, as well as the ability to model a wide range of phenomena, from predicting flooding and wildfire spread to forecasting supply chain dynamics and simulating the national electricity grid. The ability to move beyond ad hoc judgments to rigorous, data-driven decision-making while effectively managing inherent uncertainties is paramount to saving lives, protecting property, and safeguarding public health. It also enhances national security and economic strength by enabling local jurisdictions to respond to threats with nation-scale impacts.

Workshop Outcomes

Findings: The workshop identified that establishing infrastructure and operational readiness through regular exercises and well-defined scenarios represents a key centralized and high-value role for the NSCR. Participants recognized compelling success stories that demonstrate the transformative potential of applying advanced modeling and simulation capabilities to national threat response. The United States possesses significant computational and analytical assets that could be leveraged to improve decision support, including high-performance computing (HPC) infrastructure, advanced modeling capabilities, artificial intelligence (AI) and machine learning (ML) expertise, sophisticated data management platforms, and powerful visualization tools.

The workshop emphasized that continued initiatives fostering interdisciplinary collaboration among technical experts, policymakers, planners, and responders are essential to ensure both readiness and practical applicability of these capabilities. Participants identified critical investment needs in fundamental and





translational research to address gaps in uncertainty management, reliable AI-assisted decision support, fundamental limit theories, robust optimization methods, model improvement and integration, and the development of scalable workflows that effectively balance speed, accuracy, and uncertainty navigation. Clear communication strategies and visualization tools emerged as necessary components to convey complex information and uncertainties to decision makers, bridge the expert-policymaker divide, and support both immediate response needs and long-term strategic planning.

Existing Tools and Capabilities: The workshop highlighted substantial existing capabilities across multiple domains. In the energy sector, advanced geo-data tools and risk prevention models have demonstrated significant cost savings and showcased the value of AI-informed data management approaches. Environmental modeling capabilities span weather and fire prediction, earth systems modeling, and forecasting tools that provide critical insights for natural disaster response. Robust data management and visualization platforms facilitate the discovery, sharing, and utilization of data while reducing access barriers and enabling data-driven workflows.

Digital twin development efforts are advancing the integration of multiphysics models with scientific ML and uncertainty propagation for applications ranging from infrastructure monitoring to environmental prediction. Social media analytics capabilities have proven successful in enhancing situational awareness for emergency response and intelligence operations, while emerging AI data readiness assessment tools are advancing privacy-preserving collaborative model training techniques. Limit theorems are poised to provide guidance on what is possible and impossible, even when developing policies under deep uncertainty.

Potential New Uses: Workshop participants identified numerous opportunities for expanding computational decision support across critical domains. Security and infrastructure applications include enhanced screening processes that integrate multiple data sources, improved critical infrastructure resilience against both cyberattacks and natural disasters, and supply chain security models that identify vulnerabilities and assess the impacts of disruptions.

Catastrophic event response capabilities could encompass integrated modeling for extreme scenarios such as super volcanic eruptions and their long-term climate and agricultural impacts, as well as enhanced pandemic preparedness through improved modeling of information spread, biological response mechanisms, and immediate access to computational resources during outbreaks. Strategic planning applications include frameworks for decision-making under deep uncertainty and in the presence of strategic but idiosyncratic agents, which can stress-test options across various domains, as well as specialized analytics for scenarios where traditional communication and support systems are unavailable.

Citation

This report can be cited as: *Enhancing National Security and Resilience Through Computational Decision Support, DECIDE Workshop Report, Organizing Authors (Eds). Workshop Report, June 2025.*






Table of Contents

Executive Summary	2
Workshop Participants	5
Introduction	7
Aims and Goals of the DECIDE Workshop.....	8
Recommended Next Actions and High-Priority Research Directions.....	14
Conclusions	18
Acknowledgements	18
References	18
Appendices	19
Appendix A - The DECIDE Workshop Agenda	19
Appendix B - Summaries of Invited Speaker Presentations.....	22
Appendix C - Summaries of World Cafe Breakout Summaries	25

Workshop Participants

Organizing Committee

Richard Arthur, arthurr@geaerospace.com, GE Aerospace

Francis Alexander, fja@anl.gov, Argonne National Laboratory

Ray Grout, ray.grout@nrel.gov, National Renewable Energy Laboratory (NREL)

Manish Parashar, manish.parashar@utah.edu, University of Utah

Suzanne Pierce, spierce@tacc.utexas.edu, The University of Texas at Austin

Federal Points of Contact

Kirk Dohne, NITRD, NCO

Margaret Lentz, Department of Energy, Office of Science

Daniel Massey, National Science Foundation, Office of Advanced Cyberinfrastructure

Ann Schwartz, Program Manager for the National Strategic Computing Reserve (NSCR)

Government Observers

Katie Antypas, National Science Foundation, Office of Advanced Cyberinfrastructure

Hal Finkel, U.S. Department of Energy, Office of Science

Helena Fu, U.S. Department of Energy, Office of Critical and Emerging Technologies

Sheikh Ghafoor, National Science Foundation, Office of Advanced Cyberinfrastructure

Yolanda Gil, University of Southern California, Information Sciences Institute

Brian Henz, Department of Homeland Security, Science & Technology

Thuc Hoang, U.S. Department of Energy, National Nuclear Security Administration

Danielle Sumy, National Science Foundation, Technology, Innovation, and Partnerships

Shauna Sweet, DARPA, Information Innovation Office

Wen-wen Tung, National Science Foundation, Office of Advanced Cyberinfrastructure

Alonso Vera, National Aeronautics and Space Administration

Participants

Erin Acquesta, eacques@sandia.gov
Ilkay Altintas, ialtintas@ucsd.edu
Jiangzhuo Chen, chenj@virginia.edu
Clint Dawson, clint.dawson@austin.utexas.edu
Ewa Deelman, deelman@isi.edu
Lori Diachin, diachin2@llnl.gov
Emily Dietrich, edietrich@anl.gov
Kelvin Droegemeier, kkd@illinois.edu
David Ebert, ebert@ou.edu
Steve Eglash, seglash@slac.stanford.edu
Isabelle Freiling, isabelle.freiling@utah.edu
Kenny Gruchalla, kenny.gruchalla@nrel.gov
Clay Hagler, clay.hagler@microsoft.com
Beth Hartman, bethhartman@google.com
Xun Huan, xhuan@umich.edu
Kate Isaacs, kisaacs@sci.utah.edu
Kevin Jameson, kajameson@google.com
Shantenu Jha, sjha@pppl.gov
Sean Jones, sljones@anl.gov
Felice Lightstone, lightstone1@llnl.gov

Charles Macal (Chip), macal@anl.gov
Richard Macwan, Richard.Macwan@nrel.gov
Ravi Madduri, madduri@anl.gov
Jimmie McEver, jimmie.mcever@jhuapl.edu
Todd Munson, tmunson@mcs.anl.gov
Peter Nugent, penugent@lbl.gov
Jonathan Ozik, jozik@anl.gov
Valerio Pascucci, pascucci.valerio@gmail.com
Abani Patra, abani.patra@tufts.edu
Michael Rosenfield, mgrosen@us.ibm.com
Mila Rosenthal, mrosenthal@nyas.org
Sudip Seal, sealsk@ornl.gov
James Sexton, sextonjc@us.ibm.com
Aarti Singh, aarti@andrew.cmu.edu
Eric Stahlberg, eastahlberg@mdanderson.org
Dan Stanzione, dan@tacc.utexas.edu
John Towns, jtowns@illinois.edu
Nathan Urban, nurban@bnl.gov
Karen VanMeerten, vanmeert@anl.gov
Lav Varshney, varshney@illinois.edu

Workshop Synopsis

The DECIDE workshop seeks to identify opportunities and needs for developing scalable, computational decision-making methods that chart a course for impactful advancements in intelligent decision support for high-consequence decisions under uncertainty.



Introduction

The ability to provide decision support for high-consequence decisions under uncertainty is undoubtedly a prominent grand challenge of the coming decades. Rapidly evolving technologies and interconnected geopolitical systems add complexity to the decision-making process, increasing the need to incorporate advanced computational capabilities and proliferating data sources to ensure a quick response while managing uncertainty. U.S. government agencies are positioned to address a new frontier of scalable computational research, deployed through the National Strategic Computing Reserve (NSCR) [1], a strategic reserve that will provide computing/data resources and bring research expertise to bear in times of national crisis to help save lives, property, public health, and safety, or to lessen or avert the threat of a catastrophe.

The DECIDE workshop [2], held on March 4-5, 2025, brought together leading researchers, U.S. government officials, and industry experts to address the critical need for computational decision-making methods for high-consequence, complex problems. The escalating complexity and interconnectedness of modern infrastructure and ecosystems require sophisticated computational decision support capabilities for timely and informed decision-making in emergencies and high-stakes scenarios. This capability would ultimately save lives, protect property and critical infrastructure, safeguard public health, and enhance national security and economic stability. The DECIDE workshop advanced the identification of key research and engineering areas required to implement the vision of the NSCR. Specifically, the workshop was a step toward forming an interdisciplinary stakeholder community with a shared understanding of the opportunities and gaps in a national capability supporting local, front-line, and coordinated responses to emergencies.

Background

Computation and data are playing an increasingly central role in all aspects of science and society, enabling new scientific and engineering discoveries and innovations, and ensuring the nation's health, defense capabilities, and economic competitiveness. The development of scalable software and large-scale computing has heralded transformative new capabilities in modeling and simulation, even of complex systems with strategic and idiosyncratic human behavior. There exist yet unexplored opportunities to deploy effective applications of high-end computing and data systems to support time-critical decisions. One unprecedented opportunity lies within the interconnectedness and technological developments of energy, communication, transportation, domestic manufacturing, and cyberinfrastructure systems. Through ubiquitous sensing installed to control these complex systems, the data exists to model these problems before, during, and after a crisis. Solving these problems with advanced computing and data technologies within relevant time constraints can improve decision-making from ad hoc judgments to rigorous optimization problems.

This transformative power was evident through the mobilization of computing resources by the COVID-19 HPC Consortium, an innovative public-private partnership that helped address the national and global pandemic emergency. The NSCR blueprint was developed based on this experience, envisioning a national strategic reserve that could be accessed during any emergency. Like other reserve capabilities, NSCR would deepen the nation's resilience and, as appropriate, allow research expertise to be quickly brought to bear at critical junctures.





Participants in the DECIDE workshop represented an interdisciplinary stakeholder community convened to progress the NSCR vision by developing a shared understanding of the opportunities and gaps for developing, streamlining, and maintaining national capabilities to respond to urgent events.

The Imperative for Advanced Computational Decision Support

Our nation confronts a growing number of high-consequence events, including cyberattacks, pandemics, extreme weather, and infrastructure failures. In this era of deeply interconnected and complex systems, timely insight into threats and the coordination of decisions can greatly improve results across jurisdictions, as the impacts of events cascade rapidly across regions. The United States holds unparalleled computational and data capabilities across its federal agencies as well as the ability to model a tremendous variety of phenomena, from predicting flooding and wildfire spread to forecasting supply chain dynamics and simulating the national electricity grid. The ability to move beyond ad hoc judgments to rigorous, data-driven decision-making while effectively managing inherent uncertainties is paramount to saving lives, protecting property, and safeguarding public health. It also enhances national security and economic stability by enabling local jurisdictions to respond to threats with national-scale impacts.

Aims and Goals of the DECIDE Workshop

The advancement of computational and data-driven tools will be critical in addressing the challenges associated with making high-consequence decisions involving large-scale, complex systems that operate under high uncertainty and often urgency. NSCR will need these tools when called into action. The DECIDE workshop brought together leading experts to advance the understanding and implementation of these pivotal tools and to inspire transformative approaches that empower decision makers to tackle critical challenges effectively and confidently. In convening a comprehensive set of stakeholders from government, academia, and industry, the DECIDE workshop set into motion discussions and opportunities that establish, and will continue to foster, collaboration and capabilities needed to develop integrated efforts anchored in the crosscutting NSCR and related initiatives.

Workshop Focus

The workshop agenda included invited talks to frame the vision and the challenges, and brainstorming breakouts to explore key issues. The list of topics used to frame the discussion at the workshop included:

Multidisciplinary and Computational Integration for Urgent Science

- Fusing computer science, HPC, data management, and analytics, mathematics, operations research, decision sciences, and other relevant fields with problem-specific domain sciences.
 - Harnessing computing by utilizing extreme-scale HPC systems, along with broader distributed cyberinfrastructure, for scaling and integration of components.
 - Managing large ensembles of simulation and downstream analytics workflows.
- 

- Urgent computing, i.e., computing under strict time and quality constraints to support decision-making with the desired confidence within a defined time interval. How to combine multifidelity and uncertainty approaches to arrive at the best answer that can be generated in a timely manner.
- Incremental optimization: reusing previous solutions or precomputed approximate solutions to improve the performance of updating solutions based on new data.
- HPC for complex and time-sensitive simulations, scenario analysis, and stress testing.
- Streaming data management, integration, assimilation, and analysis with time and quality constraints.

Advancing Applied Mathematics in AI

- Scalable operations research: algorithms that allow routing, scheduling, and similar hard problems to be tractable through high-end (e.g., HPC) computing.
- Ingestion of high bandwidth, noisy, incomplete sensor data.
- Game theory, including mechanism design using large-scale simulation of strategic and idiosyncratic agents, multiobjective optimization, and trade-offs between optima for different objective functions.
- Stochastic multistage optimization in the face of evolving data and system states.
- Methods and metrics to implement performance guarantees into predictive ML/AI pipelines. Model guarantees, through verification and validation, are crucial to bridge the gap to a decision-making paradigm, especially for high-consequence applications.
- Leveraging AI in solving large resource allocation and logistics problems.
- AI/ML to address data quality, complexity reduction, and predictive modeling.
- Operations research for scalable stochastic optimization, resource allocation, and adaptive response.
- Information theory and decision theory foundations to understand the fundamental limits of decision-making, especially in settings of deep uncertainty

Uncertainty Management

- Developing and integrating decision-making under deep uncertainty, optimal resource allocation, computational game theory, statistics, uncertainty quantification, AI/ML, and HPC and data systems.
- Robust uncertainty management policies.
- Uncertainty quantification (UQ) to address dynamic uncertainties, model integration, and trade-off assessment.



Consensus Building, Behavior, Engagement, and Trust

- Addressing strategic behavior, fairness, accountability, transparency, interpretability, and ethics in the decision-making processes.
- Communicating uncertainty and trade-offs and enabling the incorporation of human values into the decision-making process.

Workshop Participation

The DECIDE workshop brought together leading researchers specializing in the science of decision-making and key individuals from the U.S. government and industry who rely on advancements in computational capabilities enabling decision makers to more confidently address critical challenges in managing large-scale, complex systems under high uncertainty and often urgency. By engaging scientists, policymakers, and industrial practitioners, the workshop aimed to identify opportunities and pressing needs for developing scalable, computational decision-making tools and methods that chart a course for impactful advancements in the field over the next decade.

Outcomes and Impacts

The workshop program and discussions were designed to shape the future of NSCR, enabling it to effectively address real-world, high-consequence problems. The workshop advanced innovation of pivotal tools and transformative approaches to empower decision makers to tackle critical challenges with greater clarity, confidence, and timeliness. The discussions explored integrated research and development for a practical yet rigorous framework for resource-constrained (infrastructure, computational, etc.) decision-making for multiple, large-scale, time-sensitive, high-consequence, uncertain, complex systems. The workshop promoted impactful community engagement and collaboration across government, academia, and industry to develop integrated efforts anchored in crosscutting NSCR capabilities.

Workshop Outcomes

Findings

- **There is a clear need to create and establish the NSCR** and its infrastructure and operational readiness, which will drive the development of decision support tools through regular exercises and the development of well-defined scenarios.
 - Regular exercises and the development of well-defined scenarios of emerging crises are a necessary foundation for advancing decision support systems and will have high value.
 - **There are compelling success stories that show the promise of a national capability** involving the application of advanced modeling and simulation to responding to threats.
 - **The U.S. possesses significant computational and analytical capabilities** that could be leveraged to improve decision support, including HPC, advanced modeling, fundamental theories of information,
- 



games, and decisions, AI and ML prowess, sophisticated data management platforms, and visualization tools.

- **Numerous new scenarios and applications could be developed** and drawn upon by multiple localities through rapidly deployable models and workflows.
- **Continued initiatives that foster interdisciplinary collaboration** among technical experts, policymakers, planners, and responders are essential to ensure the readiness and applicability of the capability.
- **Investment in fundamental and translational research** remains necessary to address the identified gaps in uncertainty management, reliable AI-based decision support, robust optimization methods, model improvement and integration, scalable workflows that balance speed, accuracy, and effectively navigate uncertainty, and programming systems and execution architectures for urgent computing.
- **Clear communication strategies and visualization tools** are necessary to clearly convey complex information and uncertainties to decision makers and bridge the expert-policymaker divide, inform long-term strategic planning and investments, and answer questions of importance to front-line responders.
- Developing **metrics and methods to both empower decision makers and ensure accountability** in high-stakes decision-making environments is necessary for instilling the confidence to act despite uncertainty and urgency.

Existing Tools and Capabilities

- **Risk prevention in the energy sector:** The **Offshore Reliability Modeling (ORM) Suite** developed by NETL provides groundbreaking geo-data tools and models for **offshore oil and gas risk prevention**. This has led to significant cost savings and demonstrates the value of AI and expert-informed data management.
 - **Modeling and simulation for environmental events:** Significant capabilities exist for **weather and fire prediction**, as well as **earth systems modeling**, using tools such as the DOE Energy Exascale Earth System Model (E3SM) and the NOAA Unified Forecast System.
 - **Data management and visualization:** Platforms such as the **National Data Platform (NDP)** and **Energy Data eXchange (EDX)** facilitate the management, discovery, sharing, and use of data, reducing access barriers and its use in data-driven workflows. For example, **EDX Spatial** and tools such as **EZFault3D** for subsurface geological data showcase successful visualization for expert understanding and interactive exploration.
 - **Workflow management systems (WMS):** The Pegasus WMS enables users to run their large-scale workflows across distributed and heterogeneous resources, from edge to cloud to HPC.
 - **Digital twin development:** The **DOE MusiKal Project** is working toward developing digital twin capabilities, demonstrating progress in integrating multiphysics models, scientific ML, and data. The
- 



project's focus includes successful propagation of uncertainty for global storm surge prediction and ML-based "correction" of models.

- **Social media analysis for situational awareness:** The **SMART (Social Media Analytics and Reporting Tool)** has been successfully used by the U.S. Coast Guard and intelligence fusion centers for improved situational awareness through interactive ML features.
- **Development of AI data readiness assessment tools:** Tools such as **AIDRIN** are being developed to assess the readiness of AI data for federated learning, advancing techniques for privacy-preserving collaborative model training.

Potential New Uses

- **Advanced security screening:** Employing models of **flow through airports** and integrating data from various sources to **enhance TSA screening** for new threats, considering cost-benefit trade-offs.
 - **Enhanced critical infrastructure resilience:** Applying computational support to improve the **resilience and recovery of the power grid** and other critical infrastructure, such as water and wastewater systems, against cyberattacks and natural disasters, including strategies for containment and recovery.
 - **Strengthening supply chain security:** Utilizing models to identify bottlenecks and assess the impact of disruptions in critical supply chains, such as for materials and pharmaceuticals.
 - **Response to catastrophic earth system events:** Using models such as **E3SM** and developing new integrated models to understand and plan for the long-term impacts of events like super volcano eruptions on climate, agriculture, and supply chains.
 - **Deep space exploration support:** Providing **onboard analytics and decision support** for long-duration space missions where real-time communication with ground control is limited or unavailable.
 - **Generalized decision-making under deep uncertainty:** Further developing, scaling up, and applying frameworks such as **decision-making under deep uncertainty (DMDU)** to rigorously stress-test decisions and identify robust strategies when faced with limited data and highly uncertain futures across various domains.
 - **Mathematical systems theories:** Establishing what is possible and impossible in making decisions, due to the available information and strategic behavior of agents, and thereby bound the limits of decision-making.
 - **Mechanism design through large-scale simulations of strategic and idiosyncratic agents:** Combining reinforcement learning and large language models to provide computational forms of not only wargaming but also policy design.
 - **Improved pandemic preparedness and response:** This includes leveraging models for information spread, antibody design, and ensuring immediate access to HPC resources for live modeling during outbreaks.
- 

- **Crisis-proof communications:** Advance the development of intelligent communication agents to address critical event failures that act as threat multipliers, escalating localized disruptions into national crises and undermining the capacity to coordinate essential resources and services during cascading hazard scenarios.
- **Century-scale decision intelligence:** Develop intelligent decision support systems to address planning myopia that prioritizes immediate returns over century-scale impacts. Planning frameworks to optimize long-term planning horizons that enable access to critical natural resources, maintain essential workforce expertise across key sectors, such as mining, and support national flexibility through strategic investments that enhance innovation capacity across generational timescales.

Gaps in the Knowledge Landscape and Tools

- **Data management and accessibility:** The lack of centralized, interoperable data and the need for measurements and programs to improve the AI-readiness of data.
- **Model lifecycle development and integration:** Challenges in building, updating, and curating models with new information and the replacement of outdated models.
- **Uncertainty quantification and communication:** Intuitive methods for calculating and language for communicating spatial and multimodal uncertainty in complex systems.
- **Foundational theories:** Mathematical methods to characterize the structural limits of decision-making in existing complex systems and architectural principles for new systems.
- **AI and decision-making integration:** The need for explainable AI tools and AI literacy instruction for decision makers and policymakers.
- **Workflow management systems:** The need to develop resilient WMSs that can quickly provision the necessary computational, network, and data resources and leverage them to optimize workflow execution to minimize time to solution.
- **Decision-support systems:** Limited access to testbeds for sequential decision-making and difficulties in rapid yet resilient deployment of technologies.
- **Ethical and responsibility frameworks:** Clarification of roles and responsibilities and development of best practice recommendations for AI integration.
- **Visualization and communication:** The need for collaborative visualization tools for multiple stakeholders and advancements in situated visualization.

These gaps underscore various yet interrelated challenges in integrating AI into decision-making processes, covering technical, ethical, and communication aspects.



Recommended Next Actions and High-Priority Research Directions

Models and Methods for Decision-Making

The highest priority research centers on developing computational decision-making models and methods for complex problems under uncertainty and time constraints. This research specifically targets tools, data and capabilities that could be delivered through an NSCR.

One pervasive need is to advance the practice of uncertainty management, beginning with proficiency in **uncertainty quantification (UQ)**, which involves identifying, measuring, propagating, and reducing uncertainties associated with data, models, algorithms, and predicted outcomes. Uncertainty management in UQ includes developing more robust methods to address situations where data are limited and future conditions are highly uncertain, potentially using frameworks such as **decision-making under deep uncertainty (DMDU)**, which emphasizes stress-testing decisions and exploring scenarios rather than merely prediction, as well as engineering systems theories that characterize the best that can ever be done. Crucially, time-sensitive decision-making will require balancing speed and analytical sufficiency. Research into **rapid uncertainty quantification** will be necessary to provide timely guidance that can seamlessly integrate into scalable and intuitive workflows. Furthermore, the **integration and composability of diverse models and data sources** across different scales and domains will empower decision support tools to harness multifidelity, multiscale, multidisciplinary, and multiphysics approaches.

Advancing **AI and ML** methods for decision support is another high-priority area, leveraging AI for tasks such as data analysis, predictive modeling, resource allocation, and early warnings. To overcome distrust in “black box” approaches, research is needed to ensure adequate **trustworthiness, explainability, and interpretability** to understand guidance provided by AI models used in decision support systems. Research should also focus on the **robustness of AI** for high-stakes scenarios and on addressing **ethical considerations** when integrating AI, for example, safeguards for handling sensitive information, fairness, accountability, and managing potential biases. Developing **tools and methods for sociotechnical problems** that explicitly account for human behavior, social dynamics, and include social/people components intentionally is a significant research gap and priority. Supporting these methods requires robust **data management capabilities**, including handling streaming/messy data, integrating diverse sources, improving data quality and accessibility, and tracking provenance. Additionally, reusable and easy to deploy utility services are necessary to ensure efficient data transformation and preparation for downstream analyses and modeling, leveraging workflow technologies to rapidly generate decision relevant results.

Applied Decision Science

Decision-making involves turning information into action. Decision science studies the factors that influence the decision-making process, including behavioral incentives, neuroeconomics, psychology, philosophy, information/decision theory, and game theory. These elements must be addressed when seeking to improve decision-making, particularly under the duress of high consequences, uncertainty, and urgency. Research is





needed to develop decision-making frameworks and processes that inhibit ad hoc judgments, unproductive deliberation, and avoidance of accountability. The research agenda includes developing culturally supportable methods for qualified and quantified accountability in complex, high-stakes environments, such as recording the assumptions, unknowns, and constraints that characterize the context of the decision. Research into **consensus-building and engagement models** is vital to effectively integrate multiple perspectives and knowledge resources into decision processes, particularly in multistakeholder environments. Furthermore, improving **communication and visualization methods** is a key research area to ensure complex information, uncertainties, and trade-offs are understandable and actionable for decision makers, including strategies to bridge the gap between technical experts and policymakers. The overarching goal is to foster integrated, interdisciplinary **translational research** that bridges foundational research with practical deployment and supports bidirectional learning between research and practice.

Computing and Data Cyberinfrastructure

Decision support for high-consequence, complex problems, particularly under conditions of time constraints and urgency, requires robust and reliably available cyberinfrastructure. The NSCR is envisioned as a critical national asset in emergency response, with goals including ensuring the availability of ready resources such as computing, data, software, and services that can be leveraged nimbly in times of urgent need. Achieving this necessitates investments in procuring and deploying core NSCR resources and building scale-out and redundancy into these resources to handle large-scale events. Leveraging existing HPC and other computational and data resources, including exascale systems, distributed and federated systems, and cloud services, provides significant potential for scaling and integrating computational components for effective decision-making at scale, requiring immediate access to live model simulations during crises. Research is also needed to develop robust system infrastructure and services capable of supporting automation, dynamic execution, and composable computing systems that can adapt as needs shift, while ensuring security and supply chain resilience.

Equally critical is the need for advanced data infrastructure, data management capabilities, and data access and analytics services to support computational decision support. Handling vast quantities of diverse, often messy, and streaming data is a significant challenge. High-priority research and investment areas include robust data management and integration capabilities, such as systems for streaming data management, integration, assimilation, and analysis under time and quality constraints. It is important to make relevant data collections accessible and available in compute-adjacent environments, and/or support edge-based near-data and in situ processing. Leveraging AI and expert-informed data management and architecture is necessary to improve data quality and accessibility, including specific tools for handling small and messy data, particularly for sociotechnical problems. While recognized as important, data sharing agreements were noted as outside the immediate scope of some discussions, highlighting the challenge of incentivizing open sharing and the need to understand data weaknesses and limitations. Tracking the provenance of data and models is also essential for accountability.

Complementing these computing and data needs are foundation and translational research challenges underlying urgent workflows, such as: (1) How do you drive computation through data? (2) How do you ensure





security, privacy, and trust in all aspects of the infrastructure? (3) How do you accommodate uncertainties in data and computation? (4) How do you build applications and manage workflows so that they can adapt to increase their value? (5) How do you continuously optimize workflow execution in a dynamic data-driven environment? (6) How do you develop robust system infrastructure and services to support dynamic execution? (7) How do you incorporate appropriate utility models, market models, social/trust models, etc.? and (8) How do you formulate the necessary policy and governance structures to manage operations?

The integration and interplay between computing and data infrastructure form the foundation for effective computational decision support. Research priorities include developing scalable workflows that effectively balance speed and uncertainty management, which depends on the seamless integration of computational models, tools, and frameworks with robust data capabilities. The overarching goal is to enable moving decision-making from ad hoc judgments to rigorous computational support, which requires continuous optimization in dynamic, data-driven environments. This necessitates integrated, interdisciplinary research and development programs to build a practical yet rigorous framework for resource-constrained, effective decision-making at scale. Strategic investment in the infrastructure and operational readiness of capabilities anchored by the NSCR is crucial, promoting translational research to bridge foundational work with practical deployment in communities. Ultimately, the computing and data infrastructure must support the delivery of understandable and actionable information for decision makers in time-critical situations.

Tabletop Exercises

Establishing the readiness of a computational decision support capability, particularly for high-consequence, complex problems and deployment through the NSCR, relies significantly on conducting tabletop exercises (TTXs) and drills. TTXs are crucial for ensuring this readiness by testing policies, processes, and agreements necessary for agile, effective resource mobilization in times of urgent need. Tabletop exercises provide a practical venue for testing specific computational decision support use cases. They help validate the ability to move beyond ad hoc judgments towards rigorous computational support under strict time and quality constraints, which is essential for rapid response to urgent events such as cyberattacks or pandemics. Exercises allow stakeholders to practice the rapid deployment of models and workflows required for critical scenarios, assessing if the right models and tools are available in the right locations and if a response is feasible within demanding timeframes, such as within 10 minutes versus 3 months. They are instrumental in identifying missing tools, protocols, and integration challenges, building alignment around sociotechnical systems and tools, exploring gaps and biases in decision-making processes, and incorporating diverse perspectives. Tabletop exercises are also crucial for ensuring this readiness by testing policies, processes, and agreements necessary for agile, effective resource mobilization. TTXs allow participants to practice decision-making skills in a controlled environment and are explicitly recommended for maintaining readiness and ensuring the effective use of uncertainty reduction methods.

A core investment identified is the **conducting of tabletop exercises and drills for use cases** to ensure the readiness of standing resources. To achieve this readiness, the workshop identified the need for **regular TTXs**, such as semiannual events, to continuously exercise and enhance capabilities. Regular TTXs necessitate sustained investment in the teams that prepare and conduct these exercises, including establishing **standing**





NSCR Decision Science Team(s) and Facility Team with multiyear awards. Furthermore, investment in **domain-specific expert teams** is needed for developing the realistic crisis scenarios used in the TTXs. Resources are also needed for the scenario development process itself, ensuring the crisis event, timeline, key stakeholders, facility, and necessary computational resources are identified and prepared. The development of **synthetic data** is noted as an investment that can facilitate practicing and honing decision-making skills during exercises. Crucially, investments must support the **post-exercise analysis and improvement process**, which spans several months following each TTX to assess results, identify lessons learned, and improve NSCR and domain-specific use of cyberinfrastructure.

Beyond technical validation, tabletop exercises are a vital component of fostering the interdisciplinary community necessary for effective computational decision support. They facilitate meaningful engagement and communication between technical experts, policymakers, and other stakeholders, ensuring that computational capabilities translate into scientifically sound and practically implementable decisions. Exercises are crucial for promoting the rapid uptake and understanding of complex decision support systems, exploring how these systems operate in real-world contexts involving human behavior and power dynamics. They provide a setting to test communication schemes and trade-off definitions and support continuous learning and adaptation by feeding lessons learned back into the preparedness cycle. Regular exercises contribute to the overarching goal of **translational research**, bridging foundational work with practical deployment and ensuring that computational support can deliver understandable and actionable information for decision makers in time-critical situations.

Coordination and Stewardship

The DECIDE workshop recommends the creation of a coordinating entity, for example, a consortium or a center of excellence, that can effectively advance the recommendations in this report toward the realization of NSCR. Such an entity would collaborate with various stakeholders across different sectors and agencies to develop and execute activities (e.g., tabletop exercises, challenges) and trust-building events, establish requirements, research priorities and initiatives, identify key challenges that cooperating organizations and communities would work to address, and ensure the compatibility and composability of solutions. The nature of coordination for the DECIDE events necessitates a governance structure that combines strong technical capabilities with clear information-sharing frameworks, trust-building mechanisms, and decision-making processes that balance the need for rapid responses with data quality and security requirements. An initial step towards implementation will involve defining an agile, adaptable, and shared/collaborative governance framework and organizational structure.





Conclusions

We are witnessing urgent events with increasing frequency and increasing impacts. Consequently, we must move beyond ad hoc judgments to rigorous, data-driven decision-making, while effectively managing inherent uncertainties is paramount to saving lives, protecting property, safeguarding public health, and enhancing national security and economic stability by enabling local jurisdictions to respond to threats with nation-scale impacts. The national unparalleled computational and data capabilities, services, and expertise harnessed as the NSCR can support such urgent and effective decision-making. This report summarizes the proceedings of the DECIDE workshop, where researchers, U.S. government officials, and industry experts came together to address the critical need for urgent computational decision-making methods for high-consequence, complex problems. The discussions and research direction outlined here are essential for the creation of the NSCR and realizing the NSCR vision, building an interdisciplinary stakeholder community, and forming a shared view of the opportunities and gaps for a national capability supporting local front-line response to emergencies.

Acknowledgements

Hosted by NITRD, DOE, and NSF, the DECIDE workshop sought to identify opportunities and needs for developing scalable, computational decision-making methods that chart a course for impactful advancements in the field. The organizing committee would like to thank workshop participants for their contributions, guidance, and suggestions. Special thanks are extended to Karen VanMeerten and Deborah Zemek for their essential assistance in organizing the DECIDE workshop.

References

1. National strategic computing reserve: A blueprint. <https://www.nitrd.gov/national-strategic-computing-reserve-blueprint/>, 2020.
2. DECIDE: Computational decision support for high-consequence, complex problems. <https://www.sci.utah.edu/decide>, 2025.



Appendices

Appendix A - The DECIDE Workshop Agenda

DECIDE: Computational Decision Support for High-Consequence, Complex Problems

March 4 - 5, 2025
Westin Tysons Corner
7801 Leesburg Pike
Falls Church, Virginia 22043

Tuesday, March 4, 2025

8:30 a.m. Registration and Breakfast

9:00 a.m. Welcome from NITRD

9:15 a.m. Introduction to DECIDE: Workshop Overview and Goals
• Ray Grout on behalf of the Organizing Committee

9:45 a.m. Decision-Making Motivation
*10-minute presentations with a 20-minute session for Question / Answer

- **Wildfire**
 - Ilkay Altintas (University of California, San Diego)
- **Offshore Oil**
 - Kelly Rose (National Energy Technologies Library)
- **AI and Crisis Involving Earth Systems**
 - Clint Dawson (UT Austin)
- **Center for Accelerated Real Time Analytics (CARTA)**
 - Yelena Yesha (U. Miami)

10:45 a.m. Morning Break

11:00 a.m. Decision Implication Factors on Computational Approaches
• Rick Arthur (GE Aerospace Research)

11:45 a.m. Urgent Computing
• Manish Parashar (University of Utah)

12:00 p.m. Working Lunch & Afternoon Charge

1:00 p.m. Capabilities Talks and Discussion (I)
*15-minute presentations with a 30-minute session for Question / Answer

- **Uncertainty management in decision sciences**
 - Erin Acquesta (SNL)

- **Information and Game Theory**
 - Lav Varshney (University of Illinois, Urbana-Champaign)
- **Computational Methods**
 - Alonso Vera (NASA)
- **Managing Deep Uncertainty**
 - Rob Lempert (RAND)

2:30 p.m. Afternoon Break

2:45 p.m. Broadening the Application Space: Impact of Computational Decision Support
**10 guest speakers to present 4-min. Talks, followed by 10-min. Q&A*

3:45 p.m. Breakouts with World Café Method
With relevant data collections accessible and appropriate tools, how can we pragmatically address the following:

- **Consensus Building & Engagement Models**
 - *Integration of perspectives and knowledge resources into decision making processes; how are tools for sociotechnical applications different from typical advanced computing; how can AI assist identification and rapid implementation of solutions on the ground before, during, and after events; strategies can be employed to bridge the gap between technical experts and policymakers*
- **Qualified and Quantified Accountability**
 - *Metrics to quantify accountability in complex, high stakes environments; role of transparency and public reporting in ensuring accountability*
- **Scalable Uncertainty Reduction Workflow**
 - *Methods for reducing uncertainty in decision-making processes; balancing speed and uncertainty quantification*
- **AI for Decision Sciences**
 - *Highest value uses of AI; sensitive information considerations and integrating AI into decision making processes*

5:15 p.m. Day 1 Concludes

Wednesday, March 5, 2025

8:30 a.m. Welcome and Breakfast

9:00 a.m. Report-outs: Gaps that Surfaced from World Cafe

- Consensus Building & Engagement Models
- Qualified and Quantified Accountability
- Systemic Uncertainty Reduction Workflow
- AI for Decision Sciences

10:00 a.m. Open Discussion

10:30 a.m. Morning Break

10:45 a.m. Capabilities Talks and Discussion (II)

**15-minute presentations followed by Q&A*

- **Visualization for Decision Support (30min panel + Q/A)**
 - *Kenny Gruchalla [National Renewable Energy Laboratory]*
 - *Valerio Pascucci [University of Utah]*
 - *David Ebert [Oklahoma University]*
- **Privacy Preserving Technologies**
 - *Ravi Madduri [Argonne National Laboratory]*
- **Decision Making Frameworks**
 - *Aarti Singh [Carnegie-Mellon University]*

12:30 p.m. Lunch

1:30 p.m. Breakout discussion

**Synthesis of material presented and discussion into key outcomes*

- Decision Support Needs
- Existing Tools & Capabilities
- Crosscutting RD&D Gaps
- Community Development

3:00 p.m. Breakout Session: 4x Report-out

4:00 p.m. Workshop Report Development & Action Items

5:00 p.m. Workshop Concludes

Appendix B - Summaries of Invited Speaker Presentations

Decision-Making Motivation Talks

Presentations were provided on the following areas of applications: (1) Wildfire, Dr. Ilkay Altintas, University of California, San Diego; (2) Offshore Oil, Dr. Kelly Rose, National Energy Technologies Library; (3) AI and Crisis Involving Earth Systems, Dr. Clint Dawson, The University of Texas at Austin; and (4) An overview of the Center for Accelerated Real Time Analytics (CARTA), Dr. Yelena Yesha, University of Miami. Some of the relevant points are summarized below.

Aspects of presentations focused on AI's role in managing crises involving Earth systems with an emphasis on different modeling approaches required for short-term versus long-term events and noted that although many models share similarities, their application varies based on the specific study. Dr. Dawson illustrated this with the Calf Canyon/Hermits Peak fire in New Mexico, which escalated from controlled burns due to high winds. He discussed the adaptation of the E3SM, typically used for long-term events, to regional, shorter-term scenarios, and contrasted this with the NOAA Unified Forecast System, which is better suited for short-term modeling. Another highlight was the development of digital twin frameworks that integrate data with physics-based models and AI tools. Significant challenges confront teams that aim to provide useful information for complex decision-making processes, including the decentralization of data, which complicates data gathering, and the difficulties in data assimilation and uncertainty coupling in multiphysics scenarios. Finally, speakers addressed the communication of uncertainty, with examples ranging from wildfire to hurricane forecasts and oil spills, suggesting that the current methods, such as showing 'cones of uncertainty,' may be inadequate.

Decision Implication Factors on Computational Approaches - Richard Arthur, GE Aerospace Research:

Mr. Arthur's presentation covered a range of topics to consider in computational enablement of decision-making, including risk-motivated mediocrity and waste, decision inconsistency over time and across stakeholders, informational obstacles, the importance of framing, and the emerging intersection of AI with decision science.

The presentation discussed the concept of intellectual debt, which is incurred in the use of AI while postponing investment in a rigorous understanding of how the model works (and thus the limitations of the model). He introduced the concept of "qualified accountability" and related it to the idea of a "future safety net" to lend confidence to act in the present by recording contextual sensitivities, caveats, and the ability to adapt to emergent information or changes that merit revisiting the decision at a later time.

Another important emphasis in the presentation described how problem framing can shape a continually evolving menu of alternatives in decision-making processes and highlighted an emerging field of decision science that considers cultural, psychological, and emotional factors, beyond the factual elements considered. He suggested that formalisms can be applied in this context and noted that AI is uniquely blurring the lines in this field. Finally, Mr. Arthur discussed the composition of automated reasoning systems, pointing out aspects that radically differ from conventional human reasoning, which often centers around stories rather than data.

Urgent Computing - Manish Parashar, University of Utah: Dr. Parashar outlined a vision for leveraging advanced computing as a strategic national asset in times of crisis. Drawing from the experience of the COVID-



19 HPC Consortium, it introduced the concept of the NSCR, which would ensure rapid access to computing, data, and expertise during emergencies. The talk emphasized the importance of “urgent computing,” i.e., computing under strict time and quality constraints to support real-time decision-making in scenarios such as pandemics, extreme weather, and cyberattacks. It also explored the evolving digital continuum, where edge devices, high-speed networks, and large-scale data centers enable dynamic, data-driven workflows underlying urgent decision-making. A key theme was the role of translational research in bridging foundational science with real-world deployment, and the need for robust infrastructure, governance, and policy to support the use of responsible AI to support urgent decision-making. The presentation concluded with a call to action for the research community to proactively prepare for urgent events and ensure that computing research and infrastructure can be mobilized effectively when it matters most.

Day 1 Capabilities Talks

Discussions summarize key talking points separately for (1) Uncertainty Management in Decision Sciences, (2) Information and Game Theory, (3) Computational Methods (Alonso Vera, NASA), and (4) Managing Deep Uncertainty (Robert Lempert, RAND). The following sections briefly highlight relevant points from these invited presentations:

Uncertainty Management in Decision Sciences, Erin Acquesta, Sandia National Laboratories: Dr. Acquesta’s presentation highlighted questions about translating between solution space to decision space, thinking about the time constraints and urgency around accreditation for models for use for certain problems, and segmenting the decision makers from the modeling process enough so that their time is not wasted but keeping them engaged enough to have ownership or buy-in to the models/outputs/recommendations.

Information and Game Theory, Lav Varshney, UIUC: In his presentation, Prof. Varshney discussed the foundational theories of decision-making, namely information theory and game theory. These two, and other engineering systems theories, respect the constraints of physical laws, but being theories of what can be of use, they abstract from most physical details to allow maximum freedom for ingenuity to create useful artifacts. The key result from such theories is limit theorems that establish which resources and performance criteria are fundamental and which are largely unimportant; demarcate what is possible from what is impossible, providing design insights into operating at the boundary, that is, principles/architectures for optimal designs; define fundamental benchmarks that allow an evaluation of new engineering designs on an absolute scale, rather than only compared to previous technologies; state ideals for pushing people to build technologies that approach/achieve these absolute limits; and perhaps most importantly, characterize probabilities and possibilities of what-if scenarios, especially in settings of deep uncertainty. For example, AI-driven agent-based games can be used for policy designs such as resilient food supply chains and nuclear nonproliferation.

Computational Methods - Alonso Vera, NASA: Dr. Vera discussed the role of humans not just as sensors, but also as pattern or hypothesis generators. He highlighted the unexplored potential and great need for Human-AI Interaction (HAI) in dashboards and real-time telemetry, particularly for onboard telemetrics. Dr. Vera also raised the question of how to integrate subject matter expert knowledge to work with vast datasets.





Managing Deep Uncertainty - Robert Lempert, RAND: Dr. Robert Lempert explored the concept of decision-making under conditions of deep uncertainty, emphasizing that robust decisions can and should be made even without precise predictions. Lempert argued against the traditional “predict then act” model, explaining that under deep uncertainty—when stakeholders disagree or are uncertain about future conditions and their outcomes—prediction-based decision-making can lead to brittle and inflexible plans. Instead, he advocated for a “context-first” approach, where decision makers start by stress-testing multiple potential futures through exploratory modeling and scenario analysis. Lempert highlighted the effectiveness of DMDU, showcasing methods such as scenario discovery and robust optimization that allow planners to identify resilient strategies and manage trade-offs among competing objectives. He illustrated this approach with practical examples, including COVID-19 response strategies, demonstrating how robust strategies outperformed fixed-threshold approaches in the face of unexpected developments. Ultimately, Lempert's approach encourages decision makers to trust their adaptive and robust plans rather than uncertain predictions, facilitating better policy and strategy under conditions of complexity and ambiguity.

Day 2 Capabilities Talks

Discussions summarize points separately for (5) Visualization for Decision Support (Kenny Gruchalla, National Renewable Energy Laboratory; Valerio Pascucci, University of Utah), and (6) (David Ebert, Oklahoma University); (7) Privacy Preserving Technologies (Ravi Madduri, Argonne National Laboratory); (8) Decision-making Frameworks (Aarti Singh, CMU).

Visualization for Decision Support - Kenny Gruchalla, National Renewable Energy Laboratory, and Valerio Pascucci, University of Utah: Dr. Pascucci highlighted the critical role of visualization in making large amounts of high-quality data actionable for experts. He presented a use case in precision surgery for cancer, where AI and decision support tools assist surgeons in making quick decisions about tissue removal. Pascucci differentiated between interpretability tools for ML experts and other tools for domain experts. He described a continuum of visualization methods, ranging from manual to fully automated, akin to AI-human interaction patterns. Additionally, he noted that visualization tools can function as debugging aids, enhancing the quality of recommendations.

Dr. Gruchalla emphasized the essential role of visualization in supporting decision-making within complex, data-rich environments with multiple energy systems use cases. He discussed how effective visualization can transform high-dimensional, multivariate data into actionable insights, enabling real-time risk assessment and response. Gruchalla called attention to pressing research needs in uncertainty visualization, multi-stakeholder trade-off analysis, and integrating real-time data and network structures. He highlighted emerging approaches—including immersive analytics, interactive digital twins, and context-aware visualizations—as promising pathways toward more intuitive decision-support systems. His talk underscored the importance of a multidisciplinary approach, drawing on data science, computational methods, and human-computer interaction, to address the evolving challenges of visualization for decision support.

Trustable Human-in-the-Loop in IDM for Real-World Situations - David Ebert, Oklahoma University: Dr. Ebert highlighted the role of interactive visualization in addressing context, historical information, and nondigital





data. He also emphasized the importance of balancing decision makers' skills with machine capabilities to foster effective collaboration between humans and machines.

Privacy Preserving Technologies - Ravi Madduri, Argonne National Laboratory: Dr. Ravi Madduri discussed the potential and challenges of privacy-preserving federated learning (FL), particularly emphasizing its importance for scientific and biomedical decision-making. He highlighted that AI's demand for large-scale and diverse data sets often conflicts with privacy constraints, regulatory requirements, and the sensitivity of data. Federated learning was introduced as a promising solution because it allows models to be trained collaboratively without exposing local raw data. Dr. Madduri outlined several critical algorithmic challenges, including managing communication efficiency, memory usage, energy consumption, and handling client heterogeneity. He stressed the significance of integrating differential privacy techniques to mitigate risks associated with potential data leaks from shared model updates. Furthermore, Dr. Madduri presented a vision for evolving from current unimodal applications towards robust, multimodal, continuous federated learning systems, capable of adapting in real-time to new data and changing conditions, thus enhancing the accuracy, fairness, and generalization of AI-driven decision support systems.

Decision-making Frameworks - Aarti Singh, Carnegie Mellon University: Dr. Singh emphasized the opportunities to explore complex decisions through the use and deployment of agents in human-AI partnerships and human-in-the-loop approaches. Additionally, she addressed issues around incorporating culturally feasible/acceptable considerations and developing intelligent decision support tools capable of informing long-term perspectives in planning and preparation prior to emergent events.

Appendix C - Summaries of World Cafe Breakout Summaries

Introduction and World Cafe Description

C.1 - Consensus Building and Engagement Models

C.2 - Qualified and Quantified Accountability

C.3 - Scalable Uncertainty Reduction Workflow

C.4 - AI for Decision Sciences

C.5 - General Challenges and Opportunities

Introduction to the World Cafe approach and instructions provided to the workshop participants.

Facilitator Guide for World Cafe Exercise - Opening and Convening the Group Exercise

Welcome and Introduction: “The World Cafe is a structured conversational process intended to facilitate open discussions, and link ideas within a larger group to access the collective intelligence in the room. Our goal today is to explore key topics related to high-consequence events and decision-making processes.” *(not data sharing or provisioning, but given where we are, what can we do)*





Explanation of the World Cafe Method: “The World Cafe method involves several rounds of conversation at different tables, each focused on a specific topic. Participants will rotate between tables, allowing for diverse perspectives to be shared and integrated. Each table has a host who will remain at the table to guide the discussion and capture key insights.”

Instructions for Participants:

- **Rounds of Discussion:** We will have [number] rounds of discussion, each lasting [time] minutes.
- **Rotation:** After each round, you will move to a new table. Please ensure you visit all tables.
- **Contributions:** Share your thoughts openly and listen actively to others. Every perspective is valuable.
- **Documentation:** Table hosts will capture key points and insights from each discussion.
- **Starting the Exercise:** “Let's begin our first round of discussions. Please move to your assigned tables and start exploring the topics. Enjoy the conversations!”
- Table Host Guide:
- **Welcome Participants:** Greet participants and briefly introduce the topic.
- **Facilitate Discussion:** Encourage everyone to share their thoughts on the framing questions.
- **Capture Key Insights:** Take notes on the main points and ideas discussed.
- **Summarize:** At the end of the round, summarize the key insights for the next group.

Given that relevant data collections are accessible and available in compute-adjacent environments and tools, how can we pragmatically address the following questions, while also understanding that we will be operating in a time-constrained environment?

Appendix C.1 - Summary for Consensus Building and Engagement Models

Framing Questions

- How can DECIDE ensure that multiple perspectives and knowledge resources are effectively integrated into decision-making processes for high-consequence problems? [Understanding the landscape]
 - How are tools for sociotechnical applications different from typical advanced computing, and how can AI assist or accelerate the identification and rapid implementation of solutions on the ground before, during, and after events? [Technology and tools]
 - What strategies can be employed to bridge the gap between technical experts and policymakers to ensure decisions are both scientifically sound and practically implementable? [Bridging stakeholder divides]
- 



The session emphasized the importance of ensuring multiple perspectives are integrated in all aspects of the consensus-building process, including having interdisciplinary and cross-sector teams and using multiple modes of data collection. The group also acknowledged that there are often constraints such as time, resources, and knowledge that need to be considered. The group emphasized the critical nature of communication in consensus building. They noted that inaction is itself an action that has a cost, highlighting the importance of active engagement.

The session highlighted that sociotechnical problems are different from deterministic ones, as they involve societal reactions, engagement, accountability, and values considerations. AI was suggested as potentially useful for translating, annotating, and capturing knowledge in these contexts.

Several tools were mentioned for consensus building and engagement, including scientific improv, tabletop exercises, full-scale exercises, and experimental decision-making. The group stressed the need to build a structured process for consensus building. Several approaches were discussed for bridging divides between stakeholder perspectives:

- Using large language models (LLMs) for translation
- Improving AI literacy among policymakers and experts
- Creating meaningful engagement between decision makers and technologists
- Including human and societal implications in the process
- Using “translators” to bridge language differences between groups.
- A key action item identified was the need to create an overarching structure and evaluation plan from the start of any consensus-building or engagement process.

This session underscored the complex, multifaceted nature of consensus-building and engagement, emphasizing the need for diverse perspectives, structured processes, and effective communication strategies.

Appendix C.2 - Summary for Qualified and Quantified Accountability

Framing Questions

- How can we develop metrics to quantify accountability in complex, high-stakes environments?
[Responsible AI, Liability Protections]
- What role should transparency and public reporting play in ensuring accountability in decision-making for large-scale, complex systems?

The discussion around qualified and quantified accountability metrics in complex, high-stakes environments raised questions about the scope of accountability and whose accountability should be measured. This discussion includes human analysts providing information and recommendations to decision makers, AI systems providing information, decision makers themselves, or potentially all of these parties.





In complex systems such as electric power grids, accountability becomes more challenging due to the spatial and temporal separation between causes and effects. In complex systems, it is often difficult to infer why disruptions happen, where they occur, and identifying root causes in real-world scenarios. For that reason, good explanations must back up any information or recommendations. Notably, most AI systems find it difficult to provide good explanations or causal chains in their reasoning. For AI systems, the group suggested starting with qualitative metrics. One proposed approach was using AI for “thumbs up or thumbs down” readings to build a reputation for the AI over time through use.

The importance of tracking, documenting, and recording provenance is highlighted. The group suggested that there is a need for a cultural shift in public understanding, acknowledging that knowledge evolves over time. The group recommended not only presenting outputs but also the assumptions and constraints behind them. Analytical or modeled outputs should be presented as conditional and evolving as new knowledge emerges. Additionally, the inclusion of retrospective analyses of decisions was suggested to provide lessons learned. Guidance and approaches to incorporating retrospective analyses can be gleaned from operational agencies, such as hurricane centers, already incorporate this as part of their formal process.

Regardless of what and how information or knowledge is communicated, the group wholly agreed that transparency is essential when reporting to the public and requires “good explanations.” Decision makers need to understand and be clear about the precise decision they are making. Additionally, the public must be able to understand why choices were made. In summary, the group discussion emphasized the complexity of establishing accountability metrics in high-stakes environments and underscored the critical role of transparency and clear communication in ensuring public trust and accountability.

Appendix C.3 - Summary for Scalable Uncertainty Reduction Workflow

Framing Questions

- What are the most effective methods for reducing uncertainty in decision-making processes, and how can they be implemented across sectors? [Reusable and scalable across settings]
- How can we strike a balance between the need for rapid decision-making and the need for adequate uncertainty quantification, thereby reducing uncertainty in decision-making processes? [Temporal and urgent compute]

The group identified key aspects of scalable uncertainty reduction workflows for decision-making processes. The importance of goal-oriented uncertainty quantification questions whether it is necessary to know everything about uncertainty or if it is sufficient to determine if the uncertainty is big or small. This approach suggests focusing on the most critical uncertainties relevant to the decision at hand. It is critical for analysts to first ask what information is needed from the uncertainty analysis. Once a team has identified relevant metrics of success, the group recommends using Decision-making Under Deep Uncertainty (DMDU) processes and making DMDU approaches more accessible. Rather than treating uncertainty as one big entity, the group suggests thinking about it in a structured, hierarchical manner. By adopting a structured, piecewise approach, it may be possible to automate the identification of missing data and use incremental experimental design to





determine the next most valuable piece of data to collect. Any analysis should conduct an a priori scan to determine what is actually needed to support a decision and determine whether uncertainties can actually be reduced.

There is a need to improve methods and approaches for uncertainty reduction. Conducting uncertainty analyses at scale is complex, requiring analysts to assess the level of potential consequences, devise problem frames, and identify which potential cascades in sequential processes or phenomena are significant enough for consideration, while also paying careful attention to the thresholds set in uncertainty analyses. The group raised the question of what can be done with incomplete analysis, suggesting that often decisions need to be made even when uncertainty reduction is not complete. Importantly, physics can more easily tell us what cannot happen rather than what will happen, suggesting this as a potential approach to reducing uncertainty.

Connecting uncertainty to risk depends on how well costs can be estimated and there is a need to improve methods and approaches in this area. Additionally, decision-making needs are dynamic, relying on models that frequently are developed to reflect average conditions, yet need to be adapted to new or adjacent contexts. Methods to support model adaptation and expand the capacity and the skill base for working with synthetic data collections would be useful. Additionally, to address the compound and cascading nature of many problems or the need to represent components of a decision problem across scales, modeling hierarchies need to be developed to support multiperspective and multi-scale decision-making needs.

In summary, the group encouraged the use of a multifaceted approach to scalable uncertainty reduction, emphasizing goal-oriented methods, hierarchical thinking, automation, and the importance of considering the specific context and consequences of decisions. The discussion also highlighted the need for ongoing training and readiness to effectively implement these methods. The group recommended maintaining readiness through drill exercises and training to ensure the ability to use uncertainty reduction methods effectively.

Appendix C.4 - Summary for AI for Decision Sciences

Framing Questions

- What are the highest value uses of AI to enhance the accuracy and reliability of decision-making in high-consequence scenarios? [Robustness]
- What ethical considerations must be addressed when integrating AI into decision-making processes, and how can we ensure that these considerations are managed?

AI is a powerful tool for enhancing decision sciences, particularly in data processing, anomaly detection, and facilitating communication between different stakeholders. However, the use of AI approaches also entails the need for careful consideration of sensitive information handling and the importance of clear explanations and responsibility allocation when integrating AI into decision-making processes. The group discussion highlighted several high-value applications of AI in decision sciences.

The group called for the development of best practice recommendations and protocols for using AI in decision-making processes. While AI has the capacity to challenge preconceived views, the group identified this as a





“grand challenge.” For instance, the group highlighted AI’s potential to provide early warnings of crises, enabling timely resource deployment and activation. Additionally, AI can lower barriers to entry for groups using decision support systems and act as a translator between different stakeholders involved in decision-making.

Functionally, AI approaches can streamline and accelerate data management and processing tasks. For example, AI is valuable for putting data into forms that can be easily used for decision-making and effective at finding anomalies and latent relationships in data. AI can help shrink the parameter space to support a decision-maker focusing in on aspects of a problem to diagnose real issues and options quickly.

Additionally, AI can support and streamline communications across perspectives to act as a translator between scientists, policy makers, and the public, which the group identified as a critical need. Moreover, AI can help explain and target audiences, addressing issues of lack of familiarity and different modes of understanding. AI promises to improve decision sciences practice and accelerate tasks, and it also opens several considerations regarding sensitive information that must be addressed. Most importantly, the use of AI in decision-making processes introduces risks that need to clearly identify who is responsible for the integration and define potential pitfalls or uncertainties. For example, provenance tracking and watermarking can make it evident what was used to recommend a decision. AI should explain when it is interpolating rather than extrapolating, and when AI is integrated into decision-making processes, it needs to explain why a decision is the best and how it arrived at the recommendation. Additionally, analysts using AI in applied settings need to understand the validity ranges of models, and new methods or approaches are necessary to determine how bad data can be removed without retraining the entire model. As with every element of the DECIDE workshop, user training for rapid uptake and use in urgent situations is needed to help avoid user error or misunderstanding and mitigate any bias from either AI or human users.

Appendix C.5 - Summary for Synthesis: A Session on General Challenges and Opportunities

This session provided a space for further discussion of topics that may not have been seeded in other themed discussions and became known as the “What Else?” Session. The session covered a wide range of topics, from practical considerations such as tabletop exercises and real-time AI interaction to broader issues such as conflict resolution and ethical considerations in AI-assisted decision-making. The session highlighted the universal importance of tabletop exercises to build alignment around sociotechnical systems and tools. These exercises were seen as crucial for the rapid uptake and understanding of decision support systems. The group emphasized that the nature of responses is not generic and stressed the importance of considering the specificity of decision makers. The group noted that the military is adept at decision-making processes because they think about uncertainty differently. They discussed how people have different comfort levels with what’s “good enough” or acceptable risk levels. An important point raised was that it is usually suboptimal decisions that get selected, a concept known as satisficing. The group acknowledged that what is best is not always quantifiable.

They discussed the need to think about real-world uses and power dynamics in sociotechnical systems. The group addressed the inherent conflicts in decision-making processes, noting that different decision makers have





different scales, focus, and goals. They discussed the need to accelerate the ability to resolve these conflicts. As with other sessions, the group highlighted AI's potential to act as a translator between scientists, policy makers, and the public, addressing issues of lack of familiarity and different modes of understanding. Participants also raised ethical considerations, including the need to clearly identify who is responsible in AI-assisted decision-making processes.

A key question raised was how to interact in real-time with AI systems and the group discussed the potential for using AI to provide early warnings of crises, allowing for earlier resource deployment and activation. The session highlighted the need to plan for specific cascading events, particularly when intervention is not possible. They discussed how to handle situations where the decision space becomes fundamentally different due to changing conditions. For example, the group discussed the concept of checkpoints and simulations, and how to enable adaptivity without losing time.

The session touched on the highest value uses of AI in decision sciences, including putting data into easily usable forms for decision-making, finding anomalies and latent relationships, and helping to frame issues. The fundamental need for accessible and usable data collections leads to the discussion of the need to incentivize open sharing of data. Simultaneously, there was a focus on ensuring that decision makers and technologists understand the weaknesses and limitations of their data.

