Future Directions Workshop on Causal Prediction of Human System Dynamics

May 28-29, 2025

Alex Pentland, *Stanford University / MIT*
Nina Fefferman, *University of Tennessee*
James Evans, *University of Chicago*

Prepared by:
Sithira Ratnayaka, *Virginia Tech Applied Research Corporation*
Kate Klemic, *Virginia Tech Applied Research Corporation*

Future Directions Workshop series
Workshop sponsored by the Basic Research Office, Office of the Under Secretary of Defense for Research & Engineering

VT-ARC
Virginia Tech
Applied Research Corporation

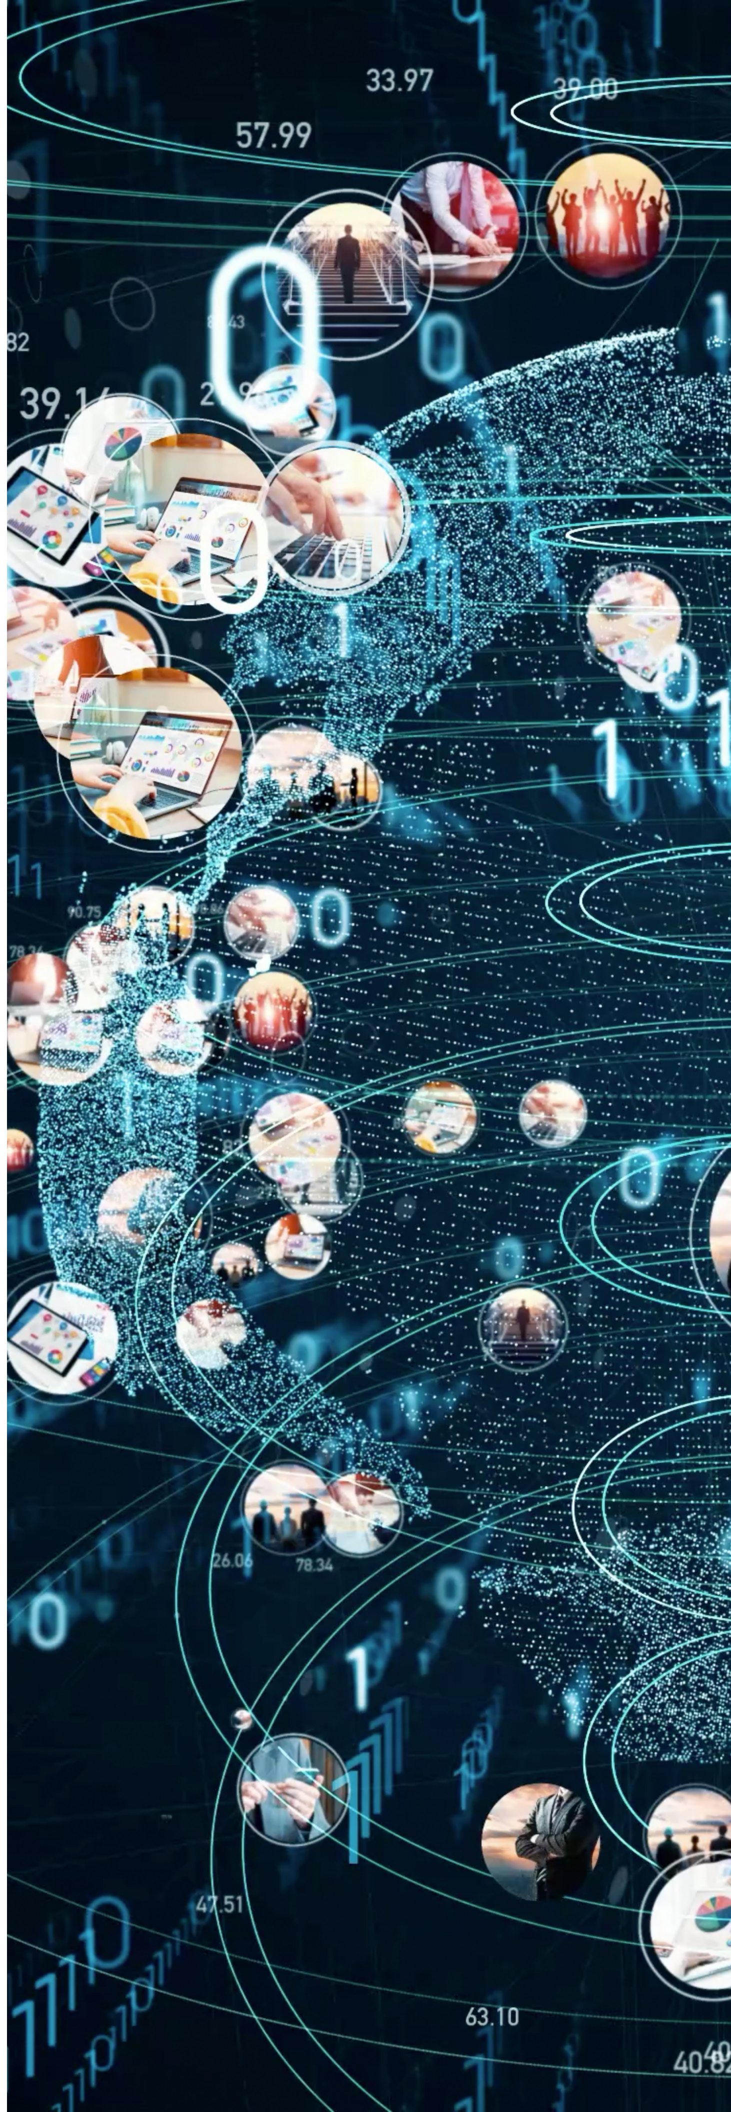

Contents

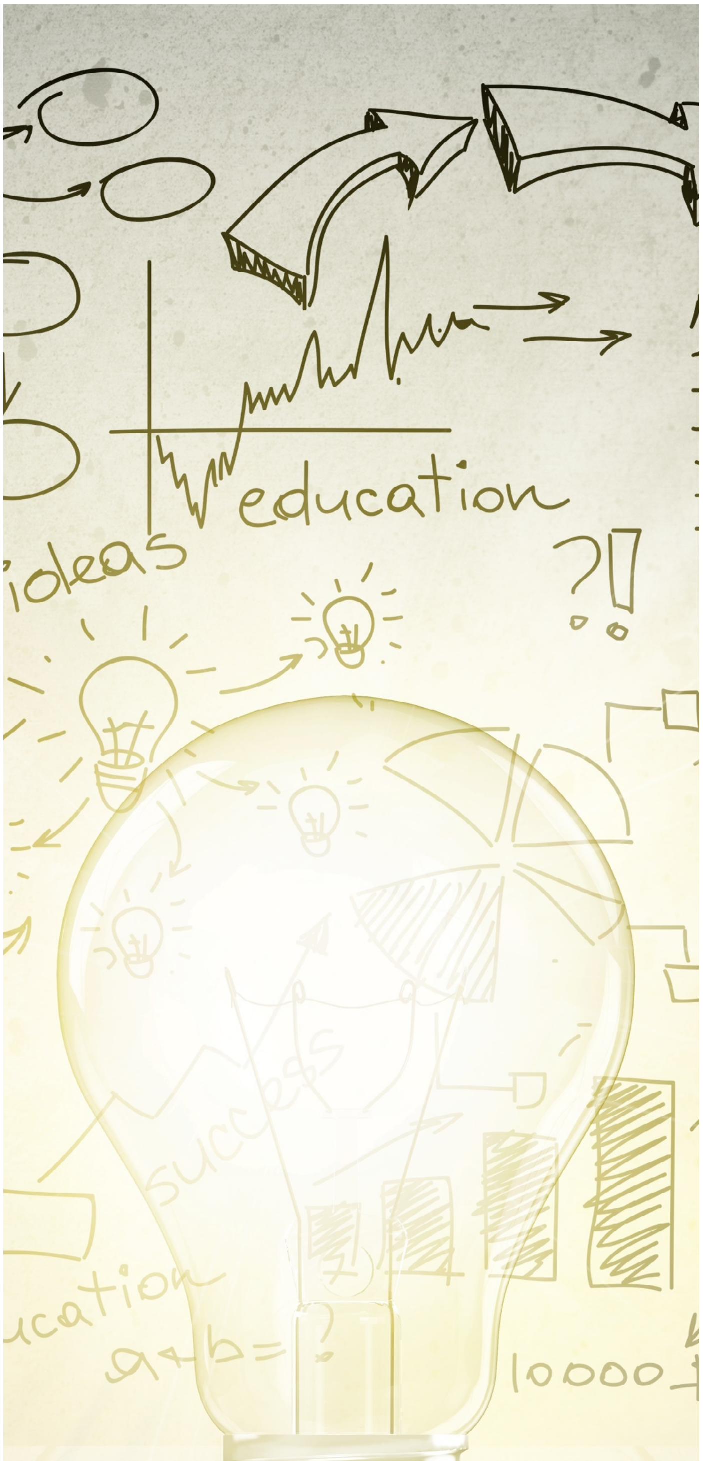

Innovation is the key to the future, but basic research is the key to future innovation.

–Jerome Isaac Friedman,
Nobel Prize Recipient (1990)

Preface

Over the past century, science and technology have brought remarkable new capabilities to all sectors of the economy; from telecommunications, energy, and electronics to medicine, transportation and defense. Technologies that were fantasy decades ago, such as the internet and mobile devices, now inform the way we live, work, and interact with our environment. Key to this technological progress is the capacity of the global basic research community to create new knowledge and to develop new insights in science, technology, and engineering. Understanding the trajectories of this fundamental research, within the context of global challenges, empowers stakeholders to identify and seize potential opportunities.

The Future Directions Workshop series, sponsored by the Basic Research Office of the Office of the Under Secretary of Defense for Research and Engineering, seeks to examine emerging research and engineering areas that are most likely to transform future technology capabilities. These workshops gather distinguished academic researchers from around the globe to engage in an interactive dialogue about the promises and challenges of each emerging basic research area and how they could impact future capabilities. Chaired by leaders in the field, these workshops encourage unfettered considerations of the prospects of fundamental science areas from the most talented minds in the research community.

Reports from the Future Direction Workshop series capture these discussions and therefore play a vital role in the discussion of basic research priorities. In each report, participants are challenged to address the following important questions:

- How will the research impact science and technology capabilities of the future?
- What is the trajectory of scientific achievement over the next few decades?
- What are the most fundamental challenges to progress?

This report is the product of a workshop held May 28-29, 2025 at Stanford University, in Palo Alto, CA, USA on the future of Causal Prediction of Human System Dynamics research, as an essential and critical aspect of future systems of modeling and integration across scales, domains, and concepts. It is intended as a resource to the S&T community including the broader federal funding community, federal laboratories, domestic industrial base, and academia.

Executive Summary

Across every domain of society – governance, health, infrastructure, defense—there is a recognition that our current methods for understanding and shaping dynamic human systems are falling short. Traditional forecasting, often grounded in static, linear, or siloed models, cannot capture the complexity of systems composed of reflexive, strategic, and socially interconnected agents. We must shift toward building modeling infrastructures that support causal reasoning, feedback-based policy design, and adaptive governance. More than gathering data, it means leveraging cutting edge methods and technologies from across disciplines to design and deploy data-aware models, that not only describe behavior but can predict it—models that operate in real time, across domains, at the scale and speed of institutional decision-making. It also demands confronting the fundamental reality that prediction itself becomes a causal force in human systems, requiring new approaches that account for the reflexive nature of social modeling.

To this effect, a science of human systems is needed that understands itself as a participant, not a bystander. This requires new conceptual formalisms, testbed infrastructure, data integration across organizational and temporal scales, and participatory architectures for experimentation. The time has come to envision a next stage of causal predictive modeling of human systems, capable of supporting dynamic coordination, detecting emerging crises, and fostering resilience across entire societies.

This report synthesizes the discussions and insights generated at the Future Directions Workshop on Causal Prediction of Human System Dynamics, convened to explore the conceptual, methodological, and infrastructural needs for modeling dynamic, adaptive, socially embedded systems. To tackle how to model, predict, and guide complex systems in which human behavior, institutional response, and environmental change co-evolve, participants identified four cross-cutting themes that arose repeatedly, converging from different starting points to arrive at similar critical focal areas of opportunity:

Feedback and Reflexivity: Participants emphasized the importance of building models that are able to anticipate and respond to how humans and institutions react to forecasts, policies, and one another.

Interoperability in Model Design: Participants called for modular, composable modeling architectures that can flexibly integrate across domains and scales. Rather than develop bespoke models for each domain or question, the community should develop shared infrastructures and conceptual grammars that allow reuse and adaptation.

Co-creation and Utility: Discussions repeatedly returned to the importance of embedding modeling practices within the institutions they aim to support—engaging stakeholders in co-development, ensuring interpretability, and recognizing models as part of a broader process of decision-making and governance.

Need for Testbeds and Experimental Platforms: Participants identified the urgent need for technologies and frameworks to support rapid, iterative, and participatory model development. These would enable the validation of models in real-world settings and provide a foundation for trust and interpretability.

The workshop discussions established three fundamental questions that challenge current approaches to human system dynamics modeling. First, how do we develop a hierarchy of evidence for causal reasoning in human systems, particularly when multiple competing theories can explain the same phenomena and traditional scientific parsimony may not apply to social complexity? Second, what happens when prediction itself becomes the primary causal force in human systems—transforming from external observation to active social intervention? Third, can we extend our temporal modeling horizons to capture longer-term dynamics while demonstrating immediate value to stakeholders? These questions highlight the field's need to build deeper epistemological and methodological foundations.

Discussion generated a clear call to action: to treat modeling not simply as a scientific endeavor, but as a civic resource—one that is transparent, participatory, and continually updated to reflect the complexity of the world it seeks to inform.

Research Opportunities

There is an immediate need to distinguish between exogenous signals and endogenous fluctuations in observational data of human systems. Immediate paths forward could include coupling fine-grained simulations with macro-level descriptors, constructing ensemble pipelines for cross-resolution translation, and experimenting with phase transition detection methods. Investments in dynamic scale-linking models and new formal tools to map feedback across levels are likely to be essential stepping stones for progress.

Conflicting goals must be able to be expressed explicitly within predictive systems. The challenge of managing polycentric governance, overlapping jurisdictional authorities, and divergent incentives among stakeholders is a central feature of real-world complexity. Tools must permit productive contestation, iterative goal refinement, and simulation of multi-actor scenario divergence.

Research Trajectory

The workshop participants developed a trajectory for the research opportunities identified for the field of embodied intelligence with a vision for the five-, ten-, and twenty-year horizons.

Five-year

The workshop participants anticipate advances at the five-year horizon that include improved models by integrating scale, enhancing social data use, and developing a facilitation agent to support group dynamics and decision-making.

- Integrate scaling dimensions (e.g., population size, time, geography) into existing research
- Define evaluation standards across domains
- Develop a process facilitation agent to:
    - Identify and articulate discrepancies
    - Detect disagreement and disengagement
    - Clarify individual and group priorities
    - Model behavioral heterogeneity
- Improve formatting, collection, and mathematical representation of social data

Ten-Year

For the ten-year horizon, the workshop participants anticipate a more advanced understanding of scaling in complex social systems through transdisciplinary integration, causal modeling across scales, and intelligent agent support, while embedding adaptive and institutional mechanisms for sustained innovation.

- Define and integrate scaling mechanisms across micro, meso, and macro levels
- Use intelligent agents for cross-domain fact finding, hypothesis testing, and detecting model misalignments
- Support co-evolution of problem definitions and modeling through human–AI feedback loops
- Address heterogeneity in social systems and promote adaptive alignment in decision processes
- Establish standards and institutions (e.g., centers of excellence, causal inference protocols) to coordinate and validate approaches

Twenty-Year

For the twenty-year horizon, the participants anticipate an enabling strategic, data-driven governance through simulation, AI, and causal modeling to support sustainable development and regulatory policy across complex social systems.

- Define intervention contexts and apply data-driven solutions (e.g., urban planning)
- Develop science-based policies for social media and autonomous systems
- Use simulation and AI to model outcomes, identify risks, and analyze economic/social dynamics
- Reframe problems using cross-domain insights and participatory AI
- Institutionalize causal frameworks for social system analysis and predictive governance

This workshop report outlines the opportunities and a path forward for research in the field of human system dynamics. A concerted effort must be made to bring together the community to address these challenges, as interdisciplinary research and collaboration by improving communication and idea-sharing within the community is imperative for the future of this field.

Three fundamental challenges distinguish this field from traditional predictive modeling: the problem of multiple valid causal hypotheses that resist Occam's Razor simplification; the reflexive nature of social systems that change in response to being predicted; and the need to balance immediate demonstrated value with longer temporal modeling horizons necessary for understanding systemic change.

Introduction

Human societies are complex systems where small shifts can trigger large-scale transformations. Economic crises, social unrest, and geopolitical realignments often emerge unexpectedly, highlighting the urgent need for tools that can explain and predict these dynamics. Despite advances in data and modeling, there is still a lack of rigorous understanding of the causal mechanisms driving sociocultural disruptions, leaving policymakers and institutions reactive rather than proactive.

One emerging phenomenon is that wars are getting longer (Figure 1). Is this because of great power competition during the cold war? Changes in military technology? Due to poor countries getting richer? Better communications technology? In order to understand these phenomena will require science progress toward understanding causal factors and how those predict human dynamics.

Figure 1: *Green, D. (2023, May 11). Why are civil wars lasting longer? From Poverty to Power.* *https://frompoverty.oxfam.org.uk/why-are-civil-wars-lasting-longer/*

New opportunities are emerging through the convergence of big data, artificial intelligence, and computational modeling. By integrating diverse datasets, from social networks and financial transactions to demographic records and anthropological observations, researchers can begin to uncover how individual actions scale into collective behaviors and identify early signals of instability.

However, realizing this vision requires overcoming significant challenges. Social systems are nonlinear and multidimensional, demanding models that can bridge scales, reconcile perspectives across disciplines, and explain cross-domain influences. At the same time, predictive capabilities raise questions about governance and ethics: understanding not just what we can predict, but how those predictions affect the systems themselves. By integrating theory, data, and computation across the natural, social, and computational sciences, we can aim to build predictive frameworks that are both scientifically rigorous and societally responsible.

We can break the overall problem of building predictive frameworks into three topics: the ability to generalize across scale and granularity of analysis, the ability to integrate predictions across domains of human activity, and the ability to conceptualize and model human dynamics within a conceptual framework that is tractable.

Scales

Bridging data across individual and population scales is essential for understanding sociocultural dynamics but remains a major challenge. While researchers can now access both granular local observations and large-scale datasets, current computational tools struggle to integrate these extremes. Agent-based models fail to scale, and population models often overlook local variability and feedback. Advancing multi-scale modeling is critical for predicting how disruptions emerge and spread through complex social systems.

Domains

Understanding sociocultural disruption requires integrating information across multiple domains such as economic, political, legal, demographic, and cultural, where human behavior simultaneously unfolds. Each domain operates with its own frameworks, languages, and models, making it difficult to exchange data and knowledge without losing critical context. Capturing these cross-domain interactions demands computational approaches that combine qualitative and quantitative representations, enabling explainable models that can characterize complex social activity more holistically.

Concepts

Interpreting complex sociocultural dynamics requires thinking across concepts to find new frameworks that balance simplification with explanatory power. While mathematical tools reduce high-dimensional data into abstractions we can visualize

"Social systems are nonlinear and multidimensional, demanding models that can bridge scales, reconcile perspectives across disciplines, and explain cross-domain influences."

and interpret, this process risks overlooking critical dynamics. Advancing our understanding demands identifying which aspects of social complexity can be safely treated as black boxes and which require formalization within unified, computationally tractable frameworks, particularly to better model and anticipate causal emergence in human systems.

Modeling efforts across many domains—epidemiology, policy, military strategy, public health, behavioral economics, and beyond—are struggling to capture the dynamic, adaptive, and socially reflexive nature of the systems they seek to represent. Across sectors, institutions are being asked to act in the face of uncertainty, feedback, and cascading systemic risk. Traditional modeling approaches, which assume that the system under observation is passive and will not change in response to the analysis itself, are insufficient. We must identify new approaches to modeling human systems that go beyond static prediction and support causal understanding, institutional decision-making, and adaptive feedback.

Improving our understanding of causality in human systems dynamics is an area of both urgency and opportunity: although decades of investment in the modeling of human systems have produced useful tools, those tools have often been brittle in the face of unexpected transitions. There are profound limits to our current toolkit. Trend extrapolation, scenario branching, and statistical prediction struggle when confronting systems with deep interdependence, rapid adaptation, and institutional reflexivity. Models that are causally grounded and operationally aligned must take priority over those that are descriptive or correlative.

"A next-generation science of causal dynamics in human systems will require fundamental rethinking of inference, scale, and the human-machine boundary."

The central objective is to move beyond asking "what will happen?" to "what makes things happen, and how might we intervene?" The current modeling landscape is both rich in isolated technical capacity and poor in integrative theory, with many contributions confined to narrow institutional, temporal, or disciplinary scopes. This lack of integration has hampered both the generation of explanatory insight and the deployment of trustworthy decision support. The field can grow beyond those limitations while avoiding overconfidence in opaque systems. Successful models must be accountable to purpose, to history, and to those impacted by decisions informed by model outputs. This accountability cannot be achieved without making assumptions explicit, ensuring interpretability, and designing for user cognition. Social and institutional dynamics should not be treated as mere noise to be abstracted away, but as critical features to be modeled with fidelity.

A next-generation science of causal dynamics in human systems will require fundamental rethinking of inference, scale, and the human-machine boundary. We must coordinate work across domains and institutions, and develop infrastructure, concepts, and talent necessary to build systems that support adaptive, ethical, and evidence-based governance in complex and evolving environments. These areas were chosen to focus the initial mapping of efforts that would have the most impact in the next 10–15 years.

In order to tackle these topics, the Future Directions Workshop for Causal Prediction of Human System Dynamics gathered researchers with a broad, diverse set of expertise, an open appreciation of different ideas and concepts, strong communication across several fields, and a far-reaching vision. New research relationships across the natural, formal, and social sciences were established and discussed, bringing together: new concepts and mathematical models to represent adaptive and interactive changes in social relations; understanding of how technological advances in data collection and modeling tools explain and affect future patterns at different sociocultural scales; and creating more holistic and transformative concepts and approaches to studying social dynamics.

The participants gathered for two days of facilitated discussion, leveraging a pre-workshop survey and framing talks by academic experts, to examine the future directions in the field of Causal Prediction of Human System Dynamics. The participants engaged in both small-group and large-group discussions around the future of Human System Dynamics. This report summarizes the discussion from the workshop relating to research challenges, research opportunities, and the ultimate trajectory to achieve the vision of Causal Prediction of Human System Dynamics.

Research Challenges

Workshop participants discussed the overarching challenges around utilization and current research around Human System Dynamics. The groups were focused around integration across scales, domains, and concepts.

Existing modeling frameworks often fail to reflect the essential properties of the systems they aim to describe. These failures arise not merely from limitations in data or computational capacity, but from deeper mismatches in conceptual design, institutional alignment, and epistemological assumptions. Addressing these failures leads to challenges that must be tackled to advance our capabilities in modeling human systems.

Scales

Models tend to either fail to adequately capture, or else misalign with, the scales at which behaviors emerge and decisions are made. Policy shifts occur at institutional scales, while public responses may unfold at interpersonal or network levels. Feedback across these levels—such as how risk perception shapes behavior, which in turn reshapes risk—are nuanced and tightly coupled with each other. This challenge is more than technical; it's epistemological. Our current modeling grammars often cannot express causal pathways that loop across time and structure.

Reflexivity and Prediction Paradoxes

A critical challenge emerges when prediction systems become sufficiently accurate and pervasive that they fundamentally restructure the social systems they aim to model. Current frameworks treat prediction as external scientific activity rather than as social intervention that changes the system under study. This reflexivity creates multiple challenges: prediction markets that commodify human behavior, power dynamics around who controls predictive systems, and the blurring of boundaries between research and deployment in connected societies. The field must grapple with whether the ultimate goal is prediction or control, and if control, who determines desirable outcomes.

"The field must grapple with whether the ultimate goal is prediction or control, and if control, who determines desirable outcomes."

Domains

There is persistent fragility in domain translation. While some modeling techniques have found success within specific sectors, few can traverse domains without reengineering their assumptions. Models that succeed in epidemiology may fail in economic policy or urban resilience because they rely too significantly on a single discipline's data, perspectives in model formulation, or frameworks for interpretation. Common pitfalls that arise from domain-specific blind spots must be avoided, such as the tendency to overprioritize quantifiable features while underrepresenting cultural or narrative dimensions. Full semantic interoperability must be developed - not only between data systems, but between the conceptual architectures that underpin model design.

Concepts

Many models operate with static rules, fixed actor types, and predefined interactions, yet in real systems, both the rules and the roles change - often in reaction to the model's own output. Feedback loops dominate many of our most important social features: wealth creation, for instance, comes not just from innovation but from attracting people to invest time and effort in innovative behavior, and reinvesting the resulting surpluses to further develop the innovation. Similarly, changes in international hostile behavior are often driven by a cycle of fear and resulting defensive behavior, a feedback loop known during the cold war as the security dilemma. Development of theories that account for feedback loops, reflexivity, endogenous change, and distributed agency will be critical to progress. This reflexivity means that models not only describe systems but participate in their evolution. Modeling practices themselves are institutional artifacts—shaped by incentives, norms, and histories that affect what questions are asked and how rigor is defined.

Models

Models must also accommodate both individual complexity (such as those arising from memory, identity, attention, and saliency) and group ideation (such as perceived norms, reputation, and top-down regulation). Existing modeling paradigms often carry implicit assumptions—about rationality, behavior, institutionality, and time—that fail under the stress of real-world complexity. Next generations of models will need to make reflexivity explicit, embrace the contingency of institutional knowledge, and incorporate social learning as a core causal process. Capturing mutual influence, nested intentionality, and structural co-determination may require new formulations for model interpretability and hierarchies of evidence. The challenges are not only technical but philosophical: how to build modeling ecosystems that treat causality as dynamic, feedback as central, and knowledge as both situated and evolving.

The future of causal modeling depends not on mastering any one level, domain, or technique—but on creating infrastructures that allow models to translate meaningfully across them.
The synthesis is not merely conceptual. It is infrastructural, organizational, and epistemic—a call for a new modeling paradigm built for a participatory, dynamic world.

Research Opportunities

The workshop participants identified the following key research opportunities to overcome the challenges:

Causal Theory for Complex Adaptive Systems: Emphasis should be placed on formalizing system dynamics that incorporate feedback, co-evolution, and uncertainty. Interdisciplinary synthesis must be accelerated to define theoretical primitives that bridge diverse modeling traditions. Novel methods should leverage insights from control theory, statistical mechanics, information theory, and learning systems to model how systems respond to intervention, evolve under constraint, and encode historical trajectory.

Tools for Structured Coordination: The design of computational platforms should support pluralistic negotiation, structured disagreement, and adaptive decision-making. These systems must enable human users to explore counterfactuals, test assumptions, and visualize tradeoffs in evolving conditions. We should build deliberative platforms that expose latent disagreement, support trust calibration, and mediate among institutions with conflicting values.

Models of Human-AI Co-evolution: Research should characterize the dynamics of interaction between human institutions and learning agents, including how goals, behaviors, and values shift in response to one another. Special attention must be given to unintended feedback, value drift, and model-induced behavior change. There is a clear need for frameworks that model dynamic adaptation between humans and algorithmic systems, identifying conditions under which reinforcement cycles produce systemic brittleness or emergent resilience.

Non-Human AI Cognition for Social Sensing: Research should explore deliberately designing AI systems with non-human sensory and cognitive capabilities specifically for detecting social patterns invisible to human perception. This includes temporal hypersensitivity for detecting micro-patterns in behavioral synchronization, network topology sensing for perceiving high-dimensional social structures, and emotional field detection for sensing collective mood 'pressure systems.' Such systems could provide regulatory insights that exceed human cognitive capacity while raising important questions about human-AI collaborative governance.

Modular Experimental Platforms: Modular, scalable “mini-labs” should be developed to facilitate controlled testing of competing modeling paradigms and decision support tools. These platforms must strike a balance between realism and tractability and should support both synthetic data generation and real-world application. Particularly, experimental designs must maintain high dimensional fidelity while enabling hypothesis-driven comparison of interpretive frames and prediction strategies.

Creative Conflict and Productive Disagreement Systems: Development of AI-mediated platforms that deliberately provoke productive disagreements between researchers and stakeholders while maintaining psychological safety. These systems should identify when different disciplinary communities use similar terminology with different meanings, detect goal misalignments across domains, and facilitate structured contestation that enhances collective problem-solving rather than creating division.

Social Phase Transitions and Early Warning Systems: Critical research is needed to identify mathematical frameworks for detecting social critical thresholds and phase transitions—moments when human systems undergo rapid, potentially catastrophic changes. This includes developing early warning signals for social system collapse, understanding when societies become unstable, and creating intervention points that can prevent cascade failures. Such work requires integrating insights from statistical mechanics, control theory, and social psychology to model how local instabilities can trigger system-wide transformations.

Investments should also be made in the development of model evaluation metrics, shared experimental benchmarks, and community infrastructure to support reproducibility, transparency, and modular reuse. In addition, the design of governance layers within models (i.e., structures that allow control over model behavior and built-in safeguards) is likely to be a fertile and necessary area of effort.

Cross-Cutting Themes

Several common threads emerged across the workshop's three organizing pillars—Scales, Domains, and Concepts—highlighting deep interconnections in the challenges and opportunities for modeling human systems:

Dynamic Feedback and Reflexivity: Many of the discussions emphasized the centrality of feedback loops and reflexivity. Whether across time horizons, institutional boundaries, or conceptual framings, participants pointed to the ways systems adapt in response to being observed, modeled, or acted upon. Reflexivity was identified as a distinctive feature of sociotechnical systems: human agents and institutions update in response to predictions, often in ways that invalidate prior expectations. This underscores the need for modeling architectures that account for these adaptive dynamics and their implications for causal inference.

Cross-Level Interactions and Emergence: Difficulty in linking micro-level behaviors to macro-level outcomes arose not only in the conversations focused on Scales, but throughout the workshop: temporal aggregation issues in scale, cross-sectoral integration challenges in domains, conceptual struggles with endogeneity and nonlinearity, etc. There was strong agreement that research should prioritize mechanisms that explain how local actions produce structural effects, and how macro-level patterns reshape micro-level incentives.

Misalignment and Translation Problems: Disciplinary silos, mismatched assumptions, and incompatible languages were identified as barriers to progress across all pillars. In scale discussions, this was manifest as failures to bridge resolutions; in domains, as policy-relevant fragmentation; and in concepts, as divergence in what different communities mean by 'causality' or 'agency.' There is a clear need to develop shared interpretive scaffolds and translation protocols that preserve contextual richness while enabling interoperability.

Endogeneity and Causal Complexity: There was broad consensus that modeling must grapple directly with endogenous change. This included recognizing that human institutions are not static, that preferences shift over time, and that structural conditions co-evolve with agent behaviors. Simple cause-effect models are insufficient, we will need to build formal tools that can represent co-determination, mutual influence, and temporally entangled causality.

Operational Accountability: We must ensure that modeling systems are not only technically robust but also accountable to their users and contexts. This includes aligning models with decision-maker goals, ensuring interpretability for stakeholders, and embedding value transparency. The need for epistemic, institutional, and ethical legitimacy will be a prerequisite for adoption and trust.

Modular, Transparent, Evolvable Tools: Future modeling platforms should be modular in construction, transparent in logic, and capable of evolving with new data and changing institutional conditions. Such systems should support structured contestation, pluralistic inputs, and iterative revision. This reflects a broader shift from predictive tools as static artifacts to prediction as a participatory, reflexive, and continually refined practice.

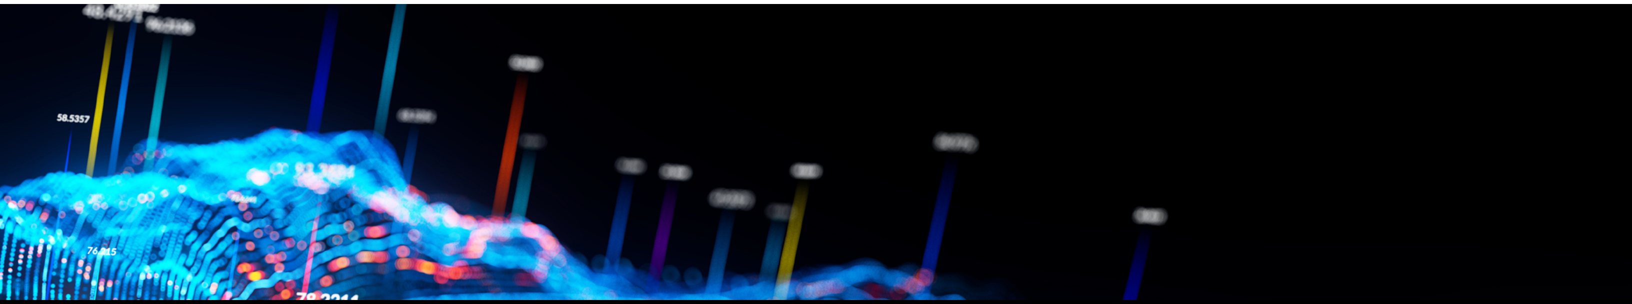

"Future modeling platforms should be modular in construction, transparent in logic, and capable of evolving with new data and changing institutional conditions."

Research Trajectory

A central aim of the workshop was to imagine a future in which causal modeling of human systems becomes more sophisticated and more impactful—integrated into the everyday functioning of democratic institutions, responsive to societal shifts, and adaptive to emergent conditions. The trajectory envisioned by participants was neither linear nor singular, but rather multi-threaded, evolving along several interdependent fronts.

Methodological: Evolution from brittle, domain-specific models toward adaptive, modular platforms that support integration across data streams, behavioral assumptions, and institutional constraints.

This includes increasing reliance on hybrid approaches that blend mechanistic simulation with machine learning, narrative modeling with sensor data, and formal inference with participatory co-design. It also includes the development of models that can remain interpretable and traceable even as they scale in complexity.

Institutional: Repositioning modeling capacity within and alongside decision-making bodies.

Participants stressed the importance of not just delivering model outputs to decision-makers, but building shared infrastructure that allows for co-creation, mutual learning, and recursive feedback. This demands investment not only in software and data infrastructure, but in new roles (e.g., model stewards, civic translators, participatory designers, etc.) who will be able to help bridge epistemic and institutional divides.

Epistemic: Developing a shared vocabulary and set of standards for modeling reflexive human systems.

This includes reconciling quantitative and qualitative traditions, building interdisciplinary curricula, and establishing norms around transparency, trust, and civic legitimacy. Building better models also means reshaping what counts as "rigor" in institutions and aligning incentives around public value, not just predictive accuracy.

Cultural and Civic: Next-generation trust in and utility from models.

Modeling must become a public resource, subject to scrutiny, shaped by diverse stakeholders, and held accountable to the systems it seeks to inform. Rather than hide behind complexity, future models should expose their assumptions, invite participation, and foster deliberation. Models must be trusted not because they are invisible, but because they are interpretable.

The trajectory of the field should not be defined by a single innovation, but by a shift in posture: from isolated analysis to embedded participation, from prediction to adaptive alignment, from control to coordination.

The 5-, 10-, and 20-year research trajectory for the cross-cutting themes are:

Dynamic Feedback and Reflexivity

In the short term, progress in dynamic feedback and reflexivity will come from building prototype models that explicitly capture how human and institutional responses reshape predictions. These efforts will rely on testbeds designed to study reflexive dynamics, such as policy feedback loops and prediction markets, while also cataloging early examples of reflexivity "failure modes" that destabilize existing systems. Within a decade, the field should be able to create hybrid models that combine insights from control theory, social learning, and adaptive AI agents, producing platforms that simulate reflexive adaptation across multiple contexts. By the twenty-year horizon, predictive ecosystems will mature into systems that can simulate reflexivity at scale, dynamically incorporating human and institutional responses into forecasts and enabling real-time governance platforms that treat reflexivity as a central design principle.

Cross-Level Interactions and Emergence

Over the next five years, researchers will begin integrating micro-level behavioral data with macro-level indicators through pilot-scale digital twins and evaluation metrics that link local behaviors to systemic outcomes. By the mid-term horizon, advances in multi-scale causal inference and agent-based modeling will allow models to generalize across diverse domains, supported by interdisciplinary centers that align theory and empirical practice. In the longer term, the field should be able to formalize "laws" of cross-scale emergence, producing models capable of predicting tipping points and systemic transformations and deploying simulations that anticipate instability and support proactive intervention across domains.

Misalignment and Translation Problems

In the near term, efforts to address misalignment and translation problems will involve identifying recurring points of divergence across disciplines - such as inconsistent terminology, metrics, and goals - and developing tools for comparing heterogeneous datasets. By the ten-year mark, the field should be able to standardize ontologies and build cross-domain infrastructures that support semantic interoperability, aided by interdisciplinary curricula and AI "translation agents" that mediate across epistemic communities. In the long term, these capabilities will evolve into global federated infrastructures that achieve seamless translation across domains, enabling real-time exchange of constructs and collaborative decision-making across disciplinary and institutional boundaries.

Endogeneity and Causal Complexity

In the short term, researchers should begin developing frameworks for modeling endogenous change, where institutions, norms, and preferences evolve in response to interventions, and piloting models that capture mutual adaptation between humans and AI systems. Within a decade, progress will focus on advancing mathematical and computational methods that can represent co-determination and temporally entangled causality, such as non-Markovian dynamics and phase transition detection, with validation protocols embedded across sectors. By the twenty-year horizon, the field should be able to establish a formal science of endogenous causality in human systems, capable of modeling adaptive institutions and structural co-evolution and deploying predictive governance systems that account for the deep interdependence of actors, rules, and environments.

Operational Accountability

In the next five years, operational accountability will be strengthened through the establishment of standards for interpretability and transparency, requiring models to document assumptions, safeguards, and stakeholder engagement. By the mid-term horizon, institutions will begin embedding stewardship roles—such as civic translators and participatory designers—within decision-making bodies, supported by audit protocols that ensure models remain aligned with public values and decision-maker goals. In the long term, predictive systems will become recognized as civic infrastructure, subject to oversight, participatory governance, and adaptive revision, where trust is built not through opacity but through transparent accountability and societal legitimacy.

Modular, Transparent, Evolvable Tools

In the near term, the field will prioritize the creation of modular platforms that allow researchers to swap and compare components such as data sources, behavioral rules, and simulation engines, supported by community repositories of reusable models. Within a decade, these platforms will mature into standardized architectures with open interfaces, enabling interoperability and iterative updating as new data, theories, and institutional contexts emerge. Over the longer term, predictive tools will evolve into dynamic ecosystems that adapt continuously to changing environments, enabling participatory and reflexive modeling practices that grow alongside the societies they aim to serve.

"The trajectory of the field should not be defined by a single innovation, but by a shift in posture: from isolated analysis to embedded participation, from prediction to adaptive alignment, from control to coordination."

Research Goals Timeline

The 5-, 10-, and 20-year research trajectory for the workshop topics are:

Five-Year	Ten-Year	Twenty-Year
	Scales	
• Integration of scales into existing lines of research • Define the dimensions of scaling across domains (population size, temporal, geography, degrees of freedom) • Develop standards of evaluation in the field • Microrandomized smartphone trials for neighborhood-level causal inference	• Define mechanisms of scaling in human social settings • Identify golden opportunities for integration across scales • Transdisciplinary integration of scaling problems • Develop and link causal methods microscopic to mesoscopic to macroscopic scales • Multi-scale social experimentation via VR/serious games • Memory-as-social-infrastructure frameworks (individual to collective memory scaling)	• Define contexts/circumstances for intervention • Enact "effective" data-driven approaches to sustainable development (e.g. Urban Planning) • Science Driven Regulation Policies for Social media and AI Agents / Autonomous systems • AI agent societies with autonomous contracts (temporal scale issues)
	Domains	
• Standardize datasets and work on privacy issues to open access to data • Coding existing exogenous variables in datasets, using datafusion techniques • Create guidelines for standardized datasets • Simulation model for offline experiments to develop/test initial predictions on simple networks • Build a coalition of the key stakeholders (government, NGOs, researchers, funders) • Developing the right research supports • Develop process facilitation agent that can identify discrepancies and articulate them for the groups • Identify underlying dimensions of heterogeneity that need to be modeled to improve fidelity of agent behavior • Identifying individual and collective priorities, relative importance of goals • Detect growing disagreement and disengagement to facilitate consensus • Blockchain-based privacy-preserving social data storage • Physics-informed neural networks for social "laws" (foundational work)	• Data governance infrastructure • Dealing with IP and privacy issues • Cooperation with agencies to do low level experiments - inject random noise to enable A/B testing • More complex simulations and hypothesis testing, generation within domain • Gathering digital trace data, recruiting participants, developing infrastructure to administer experimentsFact finder agent with cross domain knowledge • Hypothesis generation and testing engine within domains • Be able to detect discrepancies in problem definition and mental models across domains • Federated learning for sensitive cross-domain analysis • Category theory frameworks for disciplinary mismatch	• Federated learning • Decentralized storage of data that can be accessed for analysis • Curate datasets to trace causal pathways across them • Hypothesis generation, testing multiple models across domains • Cross industry, national federation to enable large scale experiments on common problems • Incentivize data sharing • Redirect discussion with insights from other domains • Simulate outcomes of different scenarios in large populations to identify potential consequences with high confidence • Suggest reframing of problem or priorities • Digital twins for entire economic systems
	Concepts	
• Correct format and collection of social data • Mathematical representations of social data • Transformer architectures for social forecasting	• C-evolution of problem definition and modelling • Feedback loop in human AI interactions • Heterogeneity in social systems • Critical thresholds and phase transition detection frameworks • Turing computability questions for AI-AI interactions (theoretical foundations)	• Robustness and stability of social systems (control theory) • Defining new form of causality of social systems and how to identify it • Deeper AI-based analysis of economic system • Mathematical resolution of computational limits for AI social systems

Example Research Areas

Advancement along this research trajectory would benefit from concerted efforts along the research themes described. Examples of these research areas include:

Next Generation Causality

Current methods of determining causality are limited to directed acyclic graphs. Thus, the assumption is that (a) all the possible causal connections are known, and (b) none of them are feedback loops. Both assumptions are clearly false in the domain of human systems, we need to extend current theory to include more general models of causality and develop statistical tools that can be efficiently and reliably used in real-world situations where data is noisy and sometimes missing.

Causal Models of Social Change

The vast majority of current work asks if any of a small number of measurable factors are important correlates of interesting outcomes. This can allow simple extrapolations to future states, but typically fails in complex, novel, or extreme circumstances. To predict transformative or disruptive events, the field must move to more dynamic models that have feedback loops and variable inter-variable coupling. Development of these models may depend on progress on the Next Generation Causality program above.

Early Detection of Extreme Events

One of the most important motivations of this area is study is early detection of extreme events. This requires adequate non-linear complex models and data from many examples of such events. Some domains, like finance, have frequent extreme events and might serve to develop initial models. Reliable early detection requires more sophisticated causal methods than we have today, and also statistical methods with the ability to deal effectively with long-tailed ("grey swan") phenomena.

Along with these three research foci, there must be sufficient financial support for research, skilled researchers, adequate compute, and large amounts of relevant data. This may mean reformation of our education and training systems, and building data and compute system adequate to the task. Programs must develop that add important, measurable value to society from the beginning, so that they can have adequate continuing support. Similarly, programs should emerge that regularly test the reliability and accuracy of methods, and that are sufficiently integrated with critical social infrastructure (e.g., healthcare, finance, etc) that deficits in performance can be quickly identifies and rectified.

Conclusions

The Future Directions Workshop on Causal Prediction of Human System Dynamics highlighted the imperative to reconceptualize the modeling of human systems in light of increasing complexity, interdependence, and reflexivity. Participants from across disciplines converged on the recognition that current modeling paradigms—often static, siloed, or limited in representational scope—are insufficient for capturing the dynamic, feedback-driven nature of social systems. Progress in this domain requires not only technical innovation, but also conceptual and institutional transformation: models must be capable of representing systems in which agents learn, adapt, and respond to being modeled.

A central insight from the workshop was the need to treat causal modeling not merely as a computational or statistical exercise, but as a scientific enterprise embedded in evolving social, institutional, and epistemic contexts. Models must account for reflexive dynamics, heterogeneity, and multi-scale interactions, while remaining interpretable, modular, and responsive to empirical feedback. Achieving this demands a shift toward hybrid approaches—integrating formal inference, simulation, participatory co-design, and narrative logic—and the development of systems and theory that allow generalization across domains without oversimplification.

The academic research community has a critical role to play in advancing this agenda. Foundational work is needed to develop new formulations of causality appropriate for systems with endogenous change and feedback. Likewise, advances in experimental infrastructure are essential for validating models under controlled yet representative conditions, particularly those involving multi-actor interaction, cross-domain effects, and institutional adaptation. The creation of modular testbeds and synthetic environments can enable structured hypothesis testing, support comparative model evaluation, and accelerate methodological innovation.

Equally important is the cultivation of a shared research culture—one that values transparency, interpretability, and interdisciplinary rigor. As modeling becomes increasingly central to decision-making in complex societal domains, researchers must engage with questions of epistemic accountability and civic relevance. This includes the development of standards for model trustworthiness, education of next-generation scholars fluent in both technical and social dimensions of modeling, and institutional support for sustained interdisciplinary collaboration. A potentially effective approach to achieve this is to integrate researchers into operational contexts so that ideas can diffuse more quickly and the science more accurately reflects reality.

Ultimately, the workshop affirmed that causal prediction of human system dynamics represents a grand challenge for the scientific community—one that sits at the intersection of theory, computation, social science, and ethics. Meeting this challenge will require a long-term, multi-faceted research effort grounded in scientific rigor and open inquiry. The path forward is not singular or prescriptive; rather, it invites diverse contributions toward a common goal: building modeling systems that are not only predictive, but explanatory, adaptive, and aligned with the evolving dynamics of the complex societies they aim to inform.

Models must account for reflexive dynamics, heterogeneity, and multi-scale interactions, while remaining interpretable, modular, and responsive to empirical feedback.

Appendix I – Workshop Attendees

Workshop Co-chairs

James Evans, *University of Chicago*
Nina H. Fefferman, *University of Tennessee, Knoxville*
Alex "Sandy" *Pentland, Stanford University*

Workshop Participants

Michael Bailey, *Social Complexity Lab at Meta*
David Barker, *American University*
Luís Bettencourt, *University of Chicago*
Damon Centola, *University of Pennsylvania*
Alin Coman, *Princeton University*
Xiaowen Dong, *University of Oxford*
Raissa D'Souza, *UC Davis*
Abhi Dubey, *Meta*
Tina Eliassi-Rad, *Northeastern University*
Fred Feinberg, *University of Michigan*
Robert Ghrist, *University of Pennsylvania*
Douglas Guilbeault, *Stanford University*
Eddie Lee, *Complexity Science Hub*
Alex Lipton, *Abu Dhabi Investment Authority*
Esteban Moro, *Northeastern University*
Matthew Silk, *University of Edinburgh*
Theodore Turocy, *The Alan Turing Institute*
Shelby Wilson, *Johns Hopkins University*
Anita Woolley, *Carnegie Mellon University*
David Wolpert, *Santa Fe Institute*
Takahiro Yabe, *New York University*

Government Observers

Jean-Luc Cambier, *Basic Research Office OUSD(R&E)*
Fariba Fahroo, *Air Force Office of Scientific Research (AFOSR)*
David Montgomery, *Basic Research Office OUSD(R&E)*
Laura Steckman, *Air Force Office of Scientific Research (AFOSR)*

VT-ARC Team

Matthew Bigman, *Virginia Tech Applied Research Corporation*
Kate Klemic, *Virginia Tech Applied Research Corporation*
Sean Lemkey, *Virginia Tech Applied Research Corporation*
Sithira Ratnayaka, *Virginia Tech Applied Research Corporation*

Workshop Participant Short Biography

Michael Bailey

Social Scientist, Meta
mikebailey@alumni.stanford.edu | https://www.linkedin.com/in/michael-bailey-a2653521/

Mike Bailey is a senior social scientist at Meta where he founded the Social Capital Lab, a research group that partners with academics and policy groups to study social networks, communities, and economic opportunity to improve society. His research group has created widely used tools for understanding social networks including the Social Connectedness Index and the Social Capital Atlas as has received widespread coverage over the years by venues such as Nature and the The New York Times.

David Barker

Professor, American University
dbarker@nsf.gov | https://www.american.edu/spa/faculty/dbarker.cfm

David C. Barker is Professor of Government at American University. He is currently on leave, serving as Director of the Social and Economic Sciences Division at the National Science Foundation. He was previously Director of the Center for Congressional and Presidential Studies (2017-2024), where he co-founded the Program on Legislative Negotiation. Earlier, he directed the the Institute for Social Research at California State University-Sacramento (2012-2017), where he founded CALSPEAKS Opinion Research. Before that, he was a professor of Political Science at University of Pittsburgh (1999-2013). Professor Barker studies American political behavior, psychology, and governance. He is the author of five books, the latest of which--Dealmakers: The Psychology of Political Compromise (Oxford University Press)--is forthcoming in early 2026.

Luís Bettencourt

Professor, University of Chicago
bettencourt@uchicago.edu | https://luisbettencourt.org

Luís M. A. Bettencourt is a Professor of Ecology and Evolution and the College at the University of Chicago. He is also Associate Faculty of the Department of Sociology and External Professor at the Santa Fe Institute. He grew up in Lisbon (Portugal) and obtained his undergraduate degree in Engineering Physics from IST Lisbon. He obtained his PhD from Imperial College London in Theoretical Physics and held postdocs and research positions at the University of Heidelberg (Germany), Los Alamos National Laboratory, MIT, and the Santa Fe Institute. His research focuses on the theory and modeling of complex systems and the processes that underlie the structure and growth of cities, in particular. He connects interdisciplinary concepts and advanced mathematics with new technologies and data to create new systems' theory and methods. This work also involves collaborations with governments, NGOs, and interdisciplinary researchers worldwide to co-produce new insights and transformative practices for sustainable development. His work is well-known academically and widely covered in the media. It has helped shape our fundamental understanding of complex systems and human societies and create novel approaches to challenges of urbanization and development.

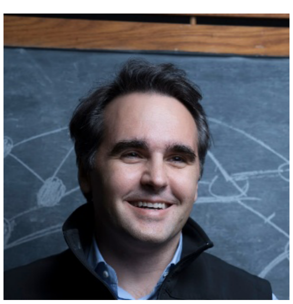

Damon Centola

Elihu Katz Professor, University of Pennsylvania
damon.centola@asc.upenn.edu | https://ndg.asc.upenn.edu/

Damon Centola is the Elihu Katz Professor of Communication, Sociology and Engineering at the University of Pennsylvania where he is Director of the Network Dynamics Group and a Senior Fellow at the Leonard Davis Institute of Health Economics. His work has received numerous awards including the Goodman Prize for Outstanding Contribution to Sociological Methodology; the James Coleman Award for Outstanding Research in Rationality and Society; and the Harrison White Award for Outstanding Scholarly Book. He was a developer of the NetLogo agent based modeling environment, and was awarded a U.S. Patent for inventing a method to promote diffusion in online networks. His work has been funded by the National Science Foundation, the Robert Wood Johnson Foundation, Facebook, the National Institutes of Health, the James S. McDonnell Foundation, the John Templeton Foundation, and the Hewlett Foundation. Popular accounts of Damon's work have appeared in The New York Times, The Washington Post, The Wall Street Journal, Wired, TIME, The Atlantic, Scientific American and CNN, among other outlets. He is a series editor for Princeton University Press, and the author of How Behavior Spreads: The Science of Complex Contagions (Princeton, 2018), and Change: How to Make Big Things Happen (Little Brown, 2021).

Alin Coman
Professor, Princeton University
Acoman@princeton.edu | https://psychology.princeton.edu/people/alin-coman

Alin Coman is a Professor of Psychology and Public Affairs, with a joint appointment between the Psychology Department and the School of Public and International Affairs. His research aims to bridge between micro-level cognitive processes and large-scale social dynamics, with a particular focus on the formation of collective memories, the dynamics of collective beliefs, and the synchronization of collective emotions.

Xiaowen Dong
Associate Professor, University of Oxford
xdong@robots.ox.ac.uk | https://eng.ox.ac.uk/people/xiaowen-dong/

Xiaowen Dong is an associate professor in the Department of Engineering Science at the University of Oxford, where he is a member of the Machine Learning Research Group. Prior to joining Oxford, he was a postdoctoral associate at the MIT Media Lab, and received his PhD degree from the École Polytechnique Fédérale de Lausanne. His main research interests concern signal processing and machine learning techniques for analysing data with complex structures, as well as their applications in social, urban, and financial network analysis.

Raissa D'Souza
Associate Dean & Professor, UC Davis
rmdsouza@ucdavis.edu | https://mae.engr.ucdavis.edu/dsouza/

Raissa D'Souza is the Associate Dean of Research for the College of Engineering at UC Davis as well as Professor of Computer Science and of Mechanical and Aerospace Engineering. She is a member of the Board of Reviewing Editors at Science, and was a Founding Lead Editor at Physical Review Research of the American Physical Society. She received a PhD in Statistical Physics from MIT in 1999, then was a postdoctoral fellow at Bell Laboratories and Microsoft Research. She is a Fellow of the American Association for the Advancement of Science (class of 2024), the American Physical Society (class of 2016), and of the Network Science Society (class of 2019). Her interdisciplinary work on network theory and complex systems spans the fields of statistical physics, theoretical computer science and applied math, and has appeared in journals such as Science, PNAS, Physical Review Letters, and Nature Physics. She has received numerous honors such as the inaugural Euler Award of the Network Science Society in 2019, the 2018 ACM Test-of-Time award, and the 2017 UC Davis College of Engineering Outstanding Mid-Career Faculty Research Award, and served as President of the Network Science Society, 2015-18.

Ahbi Dubey
Staff Research Scientist, Llama, Meta
dubeya@meta.com | https://ai.meta.com/people/1401136157273652/abhimanyu-dubey/

Abhimanyu (Abhi) Dubey is a research scientist within the Llama research group at Meta, where he is a lead researcher on visual understanding and reasoning in large language models. His most recent work is in enabling multimodal pretraining and reasoning capabilities for open-source models such as Llama 3 and Llama 4. Prior to this, he was a graduate student at the Human Dynamics group at MIT, working with Prof. Sandy Pentland on distributed online learning and differential privacy.

Tina Eliassi-Rad
Professor and Inaugural Joseph E. Aoun Chair, Northeastern University
tina@eliassi.org | http://eliassi.org

Tina Eliassi-Rad is a Professor of Computer Science and The Inaugural Joseph E. Aoun Chair at Northeastern University. She is also an external faculty member at the Santa Fe Institute and the Vermont Complex Systems Institute. Prior to joining Northeastern, Tina was an Associate Professor of Computer Science at Rutgers University; and before that a member of the technical staff at Lawrence Livermore National Laboratory. She earned her Ph.D. in Computer Sciences (with a minor in Mathematical Statistics) at the University of Wisconsin-Madison. Tina works at the intersection of AI and network science and is interested in the impact of science and technology on society. Her algorithms have been integrated into systems used by governments, industry, and open-source software. Tina received an Outstanding Mentor Award from the U.S. Department of Energy's Office of Science in 2010, became an ISI Foundation Fellow in 2019, was named one of the 100 Brilliant Women in AI Ethics in 2021, received Northeastern University's Excellence in Research and Creative Activity Award in 2022, was awarded the Lagrange Prize in 2023, and was elected Fellow of the Network Science Society in 2023.

James Evans
Professor, University of Chicago
jevans@uchicago.edu | https://sociology.uchicago.edu/directory/James-A-Evans

Max Palevsky Professor, Director, Knowledge Lab; Faculty Director, Chicago Center for Computational Social Science; External Professor, Santa Fe Institute. My research focuses on the collective system of thinking and knowing across humans and machines, including the distribution of attention and intuition, the origin of ideas and shared habits of reasoning.

Nina Fefferman
Director/Professor, University of Tennessee
nina.h.fefferman@gmail.com | https://feffermanlab.org/

Fefferman is the founding Director and PI of the US NSF Center for Analysis and Prediction of Pandemic Expansion (APPEX) and also serves as the Director of the National Institute for Modeling Biological Systems (NIMBioS). Both of these organizations are based at the University of Tennessee, Knoxville, where Fefferman is also a Professor in the Department of Ecology & Evolutionary Biology and the Department of Mathematics. Her research uses mathematical modeling to explore the behavior, evolution, and control of complex systems with application in areas from basic science (evolutionary sociobiology and epidemiology) to deployable technology (bio-security, cyber-security, and wildlife conservation). She is interested both in the application of standard modeling methods to tackle new questions, and in developing novel methods (most frequently in the area of networks science) when no current appropriate technique exists.

Fred Feinberg
Handleman Professor of Management and Statistics, University of Michigan
feinf@umich.edu | https://michiganross.umich.edu/faculty-research/faculty

Fred Feinberg is Handleman Professor of Management and Professor of Statistics, University of Michigan. He received SB degrees in Mathematics and Linguistics & Philosophy from MIT, and his PhD from the MIT-Sloan School of Management. His research builds statistical models to understand how people navigate complex environments, particularly involving sequential or contingent decisions; latent trajectories through high-dimensional spaces; and multi-agent (e.g., dyadic) choices, primarily using Bayesian, nonparametric, and dynamic programming methods. He is Departmental Editor at Production and Operations Management, former Co-Editor of Marketing Science, and served as President of ISMS. In his spare time, he likes to play classical piano and bake challah, though rarely concurrently.

Robert Ghrist
Professor and Associate Dean of Undergraduate Educ, University of Pennsylvania
ghrist@math.upenn.edu | https://www2.math.upenn.edu/~ghrist/

Robert Ghrist is the Andrea Mitchell PIK Professor of Mathematics and Electrical/Systems Engineering at Penn. He is a mathematician and academic leader serving as the Associate Dean of Undergraduate Education at Penn's School of Engineering and Applied Science. Known for his innovative research at the intersection of applied mathematics, algebraic topology, and data science, Robert specializes in leveraging topological tools -- such as sheaf theory, cohomology, and category theory -- to solve complex problems in network analysis, optimization, and information dynamics. He is a leader in integrating AI in education.

Doug Guilbeault
Assistant Professor, Stanford University
douglasguilbeault@gmail.com | https://www.gsb.stanford.edu/faculty-research/faculty/d

Guilbeault received dual bachelor's degrees in philosophy and rhetoric (with a minor in cognitive science) from the University of Waterloo, and an MA in Cognitive Linguistics from the University of British Columbia. He then completed a PhD in Communications in the Network Dynamics Group at the University of Pennsylvania's Annenberg School for Communication. He is co-director of the Berkeley-Stanford Computational Culture Lab, and a founding member of the theoretical cognitive science and machine learning collective comp-syn ("computational synesthesia"). His work has appeared in a number of top journals, including Nature, Nature Communications, The Proceedings of the National Academy of the Sciences, and Management Science, as well as in popular news outlets, such as The Atlantic, Wired, and The Harvard Business Review. He has received top research awards from The International Conference on Computational Social Science, The Cognitive Science Society, and The International Communication Association.

Eddie Lee
Postdoctoral fellow, Complexity Science Hub
edlee@csh.ac.at | https://eddielee.co

Eddie studies the role of information in the small and large living patterns around us. Examples range from the biology of neural tissue to the ecology of forests, the dynamics of armed conflict, and the processes of innovation and obsolescence in society. He is fascinated by how we paint those patterns on the shared canvas of mathematics and what the resulting similarities between the mathematical representations reveal about them. Do similarities reflect analogous function, universal dynamics, or are they (simply) artifacts of our representation? His work aims to answer these overarching questions that come together from the standpoint of information.

Alex Lipton
Global Head, Research and Development, Abu Dhabi Investment Authority
alexander.lipton@adia.ae | https://www.adia.ae/

Alexander Lipton is Global Head, Research & Development at Abu Dhabi Investment Authority, Professor of Practice at Khalifa University, Senior Founding Connection Science Fellow at MIT, and Founding Advisory Board member at ADIA Lab. Alex is a Co-Founder of Sila, and an advisory board member at several fintech companies including Swiss-Singaporean bank Sygnum and Numeraire Future Trends. In 2006-2016, Alex was Co-Head of the Global Quantitative Group and Quantitative Solutions Executive at Bank of America. Earlier, he was a senior manager at Citadel, Credit Suisse, Deutsche Bank, and Bankers Trust. Besides, Alex held visiting professorships at HUJI, EPFL, NYU, Oxford University, and Imperial College. Before becoming a quant, Alex was a Full Professor of Mathematics at the University of Illinois and a Consultant at the Los Alamos National Laboratory. Risk Magazine awarded him the Inaugural Quant of the Year Award in 2000 and the Buy-side Quant of the Year Award in 2021 (jointly with M. Lopez de Prado). Alex authored/edited thirteen books and more than a hundred scientific papers on nuclear fusion, astrophysics, applied mathematics, financial engineering, distributed ledgers, and quantum computing. He frequently gives keynote presentations at Quantitative Finance and FinTech conferences and forums worldwide.

Esteban Moro

Professor, Northeastern University

emoro@mit.edu | https://socialurban.net

Esteban Moro is a full professor and director of the Social Urban Networks (SUN) group at the Network Science Institute at Northeastern University and affiliated faculty at the MIT Media Lab. He was previously a professor and researcher at the Department of Mathematics at Universidad Carlos III de Madrid, the Sociotechnical Systems Research Center at MIT, and the University of Oxford. He holds a Ph.D. in Physics. Esteban's work lies in the intersection of big data and computational social science, with particular attention to human dynamics, collective intelligence, social networks, and urban mobility in problems like viral marketing, natural disaster management, or economic segregation in cities. He has received numerous awards for his research, and his work has appeared in major journals and is regularly covered by media outlets.

Alex Pentland

Professor, Stanford University / MIT

alexpentland@gmail.com | https://wikipedia.org/Alex_Pentland

Stanford HAI Fellow and Faculty Director DIgital Society Initative, MIT Toshiba Professor, member US Academy of Enginering, advisor UN Global Partnership Sustainable Development Data, Abu Dhabi Investment Authority Lab.

Matthew Silk

Royal Society University Research Fellow, University of Edinburgh

Matthew.Silk@ed.ac.uk | https://mattjsilk.weebly.com/

I am a Royal Society University Research Fellow in the Institute of Ecology and Evolution at the University of Edinburgh. I am broadly interested in social structure and dynamics, especially in relation to infectious disease. My work extends across human and animal social systems, applying a combination of computational and statistical modelling to better understand why social systems are structured in the way that they are, as well as the consequences of these structures for population-scale outcomes and processes. I am passionate about collaborating across disciplines, and have regularly been at the forefront of applying new developments in network science to behavioural and disease ecology. Recently, a growing focus of my research is thinking about how conflicting dynamical processes can shape the evolutionary ecology of social systems, frequently incorporating multibody interactions and distinguishing across different types of social interaction.

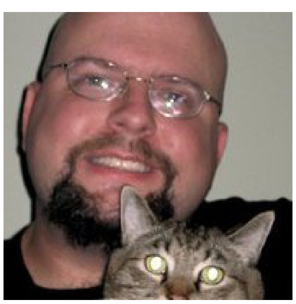

Theodore Turocy

Theme Lead, Game Theory/Collective Decision Making, The Alan Turing Institute

tturocy@turing.ac.uk | https://tturocy.github.io

I am interested in choice: specifically, how people go about making choices. My research programme extends the standard methods of economic analysis by taking account that the processes we use to make choices can affect not only what choice gets made, but how we feel about the outcomes of those choices as well. I am a theme lead at The Alan Turing Institute, where in the project Automated analysis of strategic interactions we are developing cutting-edge software tools for computing with finite games. This work is available as part of the software package Gambit: Software Tools for Game Theory. The project will support the development of automated agents able that can reason strategically about their environment, as well as enhance the reproducibility of the theoretical and empirical analysis of games. The behavioural strand of my work incorporates a richer notion of process, which takes into account the role that choice architecture, labeling, and social context, among others, colour the choice process.

Shelby Wilson
Senior Data Scientist, Johns Hopkins University Applied Physics Laboratory
Shelby.Wilson@jhuapl.edu | https://shelby-wilson.com/

Shelby Wilson is a Senior Data Scientist at the Johns Hopkins University Applied Physics Laboratory (JHU-APL). She additionally serves as the Associate Director for Education, Training, Outreach, and Participation for NSF Center of the Analysis and Prediction of Pandemic Expansion (APPEX). She is an Applied Mathematician with scientific expertise in Mathematical Biology. Her research interests include using techniques such as parameter estimation, dynamical systems, network theory, and machine learning to develop models of biological phenomena, including cancer growth, epidemiological dynamics, and social organization. Each of these interests are primarily applied to the areas of health systems operations and perspectives on interagency and whole of government interests.

David Wolpert
Professor, Santa Fe Institute
david.h.wolpert@gmail.com | https://davidwolpert.weebly.com/

David Wolpert is a professor at the Santa Fe Institute, external professor at the Complexity Science Hub in Vienna, adjunct professor at ASU, and research associate at the ICTP in Trieste. He is the author of three books (and co-editor of several more), over 250 papers, has three patents, is an associate editor at over half a dozen journals, has received numerous awards, and is a fellow of the IEEE. He has 45,000 citations, with most of his papers in thermodynamics of computation, foundations of physics, dynamics of social organizations, machine learning, game theory, and distributed optimization / control. In particular his machine learning technique of stacking was instrumental in both winning entries for the Netflix competiton, and his papers on the no free lunch theorems have over 10,000 citations. (Details at http://davidwolpert.weebly.com).

Anita Woolley
Professor, Carnegie Mellon University
awoolley@andrew.cmu.edu | https://scholars.cmu.edu/418-anita-woolley

Anita Williams Woolley is a Professor of Organizational Behavior at Carnegie Mellon University's Tepper School of Business. Dr. Woolley received her doctorate in organizational behavior at Harvard University, and her research focuses on collective intelligence in human teams and human-computer collaboration, with current projects funded by DARPA and the National Science Foundation focusing on how artificial intelligence can enhance the quality of synchronous and asynchronous collaboration in co-located and remote teams. Dr. Woolley's research has been published in Science and PNAS as well as many top journals in management, applied psychology, and computer science, and she has served as a Senior Editor at Organization Science and is a founding Associate Editor of the ACM journal Collective Intelligence.

Takahiro Yabe
Assistant Professor, New York University
ty2520@nyu.edu | https://engineering.nyu.edu/faculty/takahiro-yabe

Dr. Takahiro ("Taka") Yabe is an Assistant Professor at the Department of Technology Management and Innovation (TMI) and the Center for Urban Science + Progress (CUSP) at New York University. His research focuses on computational social science and network science approaches to model the resilience of cities to disasters, pandemics, and disruptive mobility technology, and has been published in journals including Nature Human Behaviour, PNAS, Nature Communications, and Nature Machine Intelligence.

Appendix II – Workshop Agenda and Prospectus

389 Jane Stanford Way in Stanford, CA | Simonyi Conference Center

DAY 1 – Wednesday, MAY 28, 2025

Time	Title	Speaker
8:00 – 8:15	Check-in	
8:15 - 8:20	Welcome and Introductions and Expectations	BRO & Co-chairs
8:20 -8:45	Workshop Framing Talk	Co-chairs
8:45 – 9:00	Breakout Instructions and Morning Break	
9:00 – 10:45	Working Group I: Define the Problem *Small group discussions to frame a vision for research in causal prediction of human system dynamics and identify the greatest hurdles to achieving it.* Group A – Scales Group B – Domains Group C – Concepts	
10:45 – 11:00	BREAK - *Transition to main conference room and leads prepare outbriefing*	
11:00 –12:00	Working Group 1: Outbriefing	
12:00 – 1:00	LUNCH	
1:00 – 3:45	Working Group II: Technical Capabilities and Opportunities Group A – Scales Group B – Domains Group C – Concepts	
3:45 – 4:00	BREAK - *Transition to main room and leads prepare outbriefing*	
4:00 – 4:45	Working Group II: Outbriefing	
4:45 – 5:00	Summary of Day	Co-chair(s)
5:00	MEETING ADJOURNED FOR THE DAY	

Office of the Under Secretary of Defense for
Research and Engineering OUSD(R&E)

Future Directions Workshop: Causal Prediction of Human System Dynamics

May 28–29, 2025
Palo Alto, CA

DAY 2 – Thursday, MAY 29, 2025

Time	Title	Speaker
8:00 – 8:15	Check-in	
8:15 – 8:30	Welcome	
8:30 – 9:00	Day 1 Recap	Co-chairs
9:00 – 10:00	**What's Missing?** Discussion of topics which did not fit into the framework of day 1 but need to be discussed.	
10:00 – 10:15	BREAK	
10:15 – 11:30	**Big Questions** Discussion of particularly far-out (or long-term), high-risk, high-impact ideas.	
11:30 – 11:50	Discussion of Key Ideas / Components for Report	
11:50 – 12:00	Closing Remarks Co-chairs	
12:00	Meeting Adjourned	

Office of the Under Secretary of Defense for Research and Engineering OUSD(R&E)

Future Directions Workshop: Causal Prediction of Human System Dynamics

May 28–29, 2025
Palo Alto, CA

Future Directions Workshop
Causal Prediction of Human System Dynamics

Basic Research Office, Office of the Under Secretary of Defense (R&E)

28-29 May 2025

Stanford, California

Co-chairs: Alex Pentland (Stanford/MIT), Nina Fefferman (University of Tennessee), James Evans (University of Chicago)

Today we cannot reliably predict the dynamics of sociocultural systems, even to say whether an emergent dynamic is locally stable or leading to imminent large-scale shifts, such as societal collapse, financial crises, wars, panics, etc. Further, we don't have rigorous, measurable, reliable means for understanding the causes of these sociocultural disruptions and so cannot reliably work to prevent them. In recent years, however, research leveraging multidimensional data covering large populations has demonstrated that it may be possible to quantify and condition existing social science hypotheses concerning causal factors, and to understand how these factors interact over relatively long time frames.

Studies to date have been limited in scale, domain, and conceptual breadth. Consequently, no coherent, quantitative, and unified view of sociocultural disruption has been possible. The science has to date been limited by issues of extreme ranges in variability, scale, context, change, non-linearity, and generalizability. To advance the cutting edge of the field, we must discover new ways to both consolidate/refine existing approaches and generate novel methods to better understand the causal factors that govern social complexity at different scales, e.g., temporal, geographic, and structural, and across cultural, political, economic, and other, similarly diverse contexts.

Recent advances in big data, analytics, and artificial intelligence now not only enable interrogation into a broader body of social data but challenge the field to develop novel theories and methods to understand, characterize, and predict social dynamics. Nevertheless, important challenges remain surrounding data quality and access, outcome measurement, and causality that must be answered in order to further advance the field of computational social science toward truly revolutionary capabilities with rigorously derived accuracy, quantitative predictability, and sociocultural explainability, including of variation(s). The assessment of long-term research trajectories that could potentially achieve these goals forms the basis of this Future Directions Workshop.

Fundamental questions of critical interest include, in particular, the following themes:

Theme 1: Integration across scales. Large amounts of data can be accessed from social networks, and aggregated data describing the dynamics of large segments of the population have also become available (e.g. transactional data). At the other end, we have pointwise, anthropological or psychological data, obtained by sampling locally via embedded researchers. Often analyses focus on the production of disruption, while ignoring the environments in which disruption propagates and evolves. The connection between extreme scales is particularly difficult to establish in a rigorous manner. Existing computational tools do not provide means of bridging them in a rigorous way. Agent-based modeling, for example, does not scale through multiple orders of magnitude in number of agents; multi-scale aggregation of agents and their dynamics is not an established field that can yet be relied upon. Large-scale population model formulations do not easily account for variability at smaller scales or bi-directional coupling between scales. Combined with significant gaps in real data mentioned above, challenges in building comprehensive multi-scale modeling approaches and tools are formidable. What are the prospects for achieving this goal, and what are the advances in novelty detection and emergence that might contribute to it?

Theme 2: Integration across domains. Human behavior occurs simultaneously in multiple domains, emitting and relying upon multiple signals, which makes the understanding and prediction of disruption an intrinsically high-dimensional representation problem. Social science approximately classifies these dimensions according to broad institutional and structural frameworks: economic, political, legal, demographic and cultural variables. This leads to the fundamental question of how to describe exchanges between these domains. These can include the mental frameworks of respective agents, the "language" used to describe the states and models developed for various domains, and the polysemy of robust actions that play across domains. And what are the appropriate bases for exchanging data and knowledge between domains? What information can be transmitted across domains, and what is lost? The accurate characterization of complex social activity calls for both qualitative and quantitative data representations, domain-dependent features, explainable knowledge, as well as corresponding approaches to operate on these representations, via computational models.

Theme 3: Integration across concepts. We must also face the problem of interpretation. Because of the complexity of these phenomena, we rely on mathematical tools to extract the "relevant" information and reduce everything to low-dimensional abstractions that we can visualize, imagine, and interpret. Simplification is necessarily limiting. What new, fundamental concepts underlying socio-cultural complexity would help advance our comprehension? How can we achieve deeper understanding, and what are practical and fundamental limitations in our ability to simplify social complexity? What can be safely treated as a black-box, and what needs to be formalized with a unifying language, mapped onto mathematical terms and algorithmic processes? Causal emergence is a characteristic of highly complex systems: what can be said about our ability to model and anticipate it?

It is critical to be aware that new tools for understanding and predicting sociocultural disruption will inevitably be used to modify the behavior itself. As a standard example, if one were to develop a more accurate prediction tool for financial stock valuation, this tool would rapidly be deployed to modify the stock market itself. Consequently, governance norms—such as how to communicate "proper" use, limitations, warnings, and ethical responsibilities are vital to develop along with the social science. We must accurately characterize domains of use and reliability so we can know the confidence level we can place in such modeling and prediction.

This workshop calls for a broad, diverse set of expertise, an open appreciation of different ideas and concepts, strong communication across several fields, and a far-reaching vision. New research relationships across the natural, formal, and social sciences must be established and discussed, bringing together: new concepts and mathematical models to represent adaptive and interactive changes in social relations; understanding of how technological advances in data collection and modeling tools explain and affect future patterns at different sociocultural scales; and creating more holistic and transformative concepts and approaches to studying social dynamics.