

YOLANDA GIL
University of Southern California

SUZANNE A. PIERCE
The University of Texas Austin

HASSAN BABAIE
Georgia State University

ARINDAM BANERJEE
University of Minnesota

KIRK BORNE
Booz Allen Hamilton

GARY BUST
Johns Hopkins University

MICHELLE CHEATHAM
Wright State University

IMME EBERT-UPHOFF
Colorado State University

CARLA GOMES
Cornell University

MARY HILL
University of Kansas

JOHN HOREL
University of Utah

LESLIE HSU
Columbia University

JIM KINTER
George Mason University

CRAIG KNOBLOCK
University of Southern California

DAVID KRUM
University of Southern California

VIPIN KUMAR
University of Minnesota

PIERRE LERMUSIAUX
Massachusetts Institute of Technology

YAN LIU
University of Southern California

CHRIS NORTH
Virginia Tech

VICTOR PANKRATIUS
Massachusetts Institute of Technology

SHANAN PETERS
University of Wisconsin-Madison

BETH PLALE
Indiana University Bloomington

ALLEN POPE
University of Colorado Boulder

SAI RAVELA
Massachusetts Institute of Technology

JUAN RESTREPO
Oregon State University

AARON RIDLEY
University of Michigan

HANAN SAMET
University of Maryland

SHASHI SHEKHAR
University of Minnesota

KATIE SKINNER
University of Michigan

PADHRAIC SMYTH
University of California Irvine

BASIL TIKOFF
University of Wisconsin-Madison

LYNN YARMEY
National Snow and Ice Data Center

JIA ZHANG
Carnegie Mellon University

A research agenda for intelligent systems that will result in fundamental new capabilities for understanding the Earth system.

Intelligent Systems for Geosciences: An Essential Research Agenda

MANY ASPECTS OF geosciences pose novel problems for intelligent systems research. Geoscience data is challenging because it tends to be uncertain, intermittent, sparse, multiresolution, and multi-scale. Geosciences processes and objects often have amorphous spatiotemporal boundaries. The lack of ground truth makes model evaluation, testing, and comparison difficult. Overcoming these challenges requires breakthroughs that would significantly transform intelligent systems, while greatly benefitting the geosciences in turn. Although there have been significant and beneficial interactions between the intelligent systems and geosciences communities,^{4,12} the potential for synergistic research in intelligent



systems for geosciences is largely untapped. A recently launched Research Coordination Network on Intelligent Systems for Geosciences followed a workshop at the National Science Foundation on this topic.¹ This expanding network builds on the momentum of the NSF EarthCube initiative for geosciences, and is driven by practical problems in Earth, ocean, atmospheric, polar, and geospace sciences.¹¹ Based on discussions and activities within this network, this article presents a research agenda for intelligent systems inspired by geosciences challenges.

Geosciences research aims to understand the Earth as a system of complex highly interactive natural processes and their interactions with human activities. Current approaches have fundamental shortcomings given the complexity of geosciences data. First, using data alone is insufficient to create models of the very complex phenomena under study so prior theories need to be taken into account. Second, data collection can be most effective if steered using knowledge about existing models to focus on data that will make a difference. Third, to combine disparate data and models across disciplines requires capturing and reasoning about extensive qualifications and context to enable their integration. These are all illustrations of the need for knowledge-rich intelligent systems that incorporate significant amounts of geosciences knowledge.

The article begins with an overview of research challenges in geosciences. It then presents a research agenda and vision for intelligent system to address those challenges. It concludes with an overview of ongoing activities in the newly formed research network of intelligent systems for geosciences that is fostering a community to pursue this interdisciplinary research agenda.

The pace of geosciences investigations today can hardly keep up with the urgency presented by societal needs to manage natural resources, respond to geohazards, and understand the long-term effects of human activities on the planet.⁶⁻¹¹ In addition, recent unprecedented increases in data availability together with a stronger emphasis on societal drivers emphasize the need for research that crosses over traditional

knowledge boundaries. Different disciplines in geosciences are facing these challenges from different motivations and perspectives:

► **Forecasting rates of sea level change in polar ice shelves:** Polar scientists, along with atmospheric and ocean scientists, face an urgent need to understand sea level rise around the globe. Ice-shelf environments represent extreme environments for sampling and sensing. Current efforts to collect sensed data are limited and use tethered robots with traditional sampling frequency and collection limitations. The ability to collect extensive data about conditions at or near the ice shelves will inform our understanding about changes in ocean circulation patterns, as well as feedbacks with wind circulation. New research on intelligent sensors would support selective data collection, on-board data analysis, and adaptive sensor steering. New submersible robotic platforms could detect and respond to interesting situations while adjusting sensing frequencies that could be triggered depending on the data being collected in real time.

► **Unlock deep Earth time:** Earth scientists focus on understanding the dynamics of the Earth, including the interior of the Earth or *deep Earth* (such as tectonics, seismology, magnetic or gravity fields, and volcanic activity) and the near-surface Earth (such as the hydrologic cycle, the carbon cycle, the food production cycle, and the energy cycle). While collecting data from the field is done by individuals in select

locations, the problems under consideration cover spatially vast regions of the planet. Moreover, scientists have been collecting data at different times in different places and reporting results in separate repositories and often unconnected publications. This has resulted in a poorly connected collection of information that makes wide-area analyses extremely difficult and is impossible to reproduce. Earth systems are integrated, but current geoscience data and models are not. To unravel significant questions about topics, such as Deep Earth Time, geoscientists need intelligent systems to efficiently integrate data from disparate locations, data types, and collection efforts within a wide area.

► **Predict critical atmosphere and geospace events:** Atmospheric and geospace science research aims to improve understanding of the Earth's atmosphere and its interdependencies with all of the other Earth components, and to understand the important physical dynamics, relationships, and coupling between the incident solar wind stream, and the magnetosphere, ionosphere, and thermosphere of the Earth. Atmospheric research investigates phenomena operating from planetary to micro spatial scales and from millennia to microseconds. Although the data collected is very large, it is miniscule given the complexity of the phenomena under study. Therefore, the data available must be augmented with knowledge about physical laws underlying the phenomena in order to generate effective models.

► **Detect ocean-land-atmosphere-ice interactions:** Our ability to understand the Earth system is heavily dependent on our ability to integrate geoscience models across time, space, and discipline. This requires sophisticated approaches that support composition and discover structure, diagnose, and compensate for compound model errors and uncertainties, and generate rich visualizations of multidimensional information that take into account a scientist's context.

The accompanying figure illustrates intelligent systems research directions inspired by these geoscience challenges, organized at various scales. Studying the Earth as a system requires fundamentally new capabilities to collect

» key insights

- **Advances in artificial intelligence are needed to collect data where and when it matters, to integrate isolated observations into broader studies, to create models in the absence of comprehensive data, and to synthesize models from multiple disciplines and scales.**
- **Intelligent systems need to incorporate extensive knowledge about the physical, geological, chemical, biological, ecological, and anthropomorphic factors that affect the Earth system while leveraging recent advances in data-driven research.**
- **A new generation of knowledge-rich intelligent systems have the potential to significantly transform geosciences research practices.**

data where and when it matters, to integrate isolated observations into broader studies, to create models in the absence of comprehensive data, and to synthesize models from multiple disciplines and scales. Advances in intelligent systems to develop more robust sensor platforms, more effective information integration, more capable machine learning algorithms, and intelligent interactive environments have the potential to significantly transform geosciences research practices and expand the nature of the problems under study.

A Roadmap for Intelligent Systems Research with Benefits to Geosciences

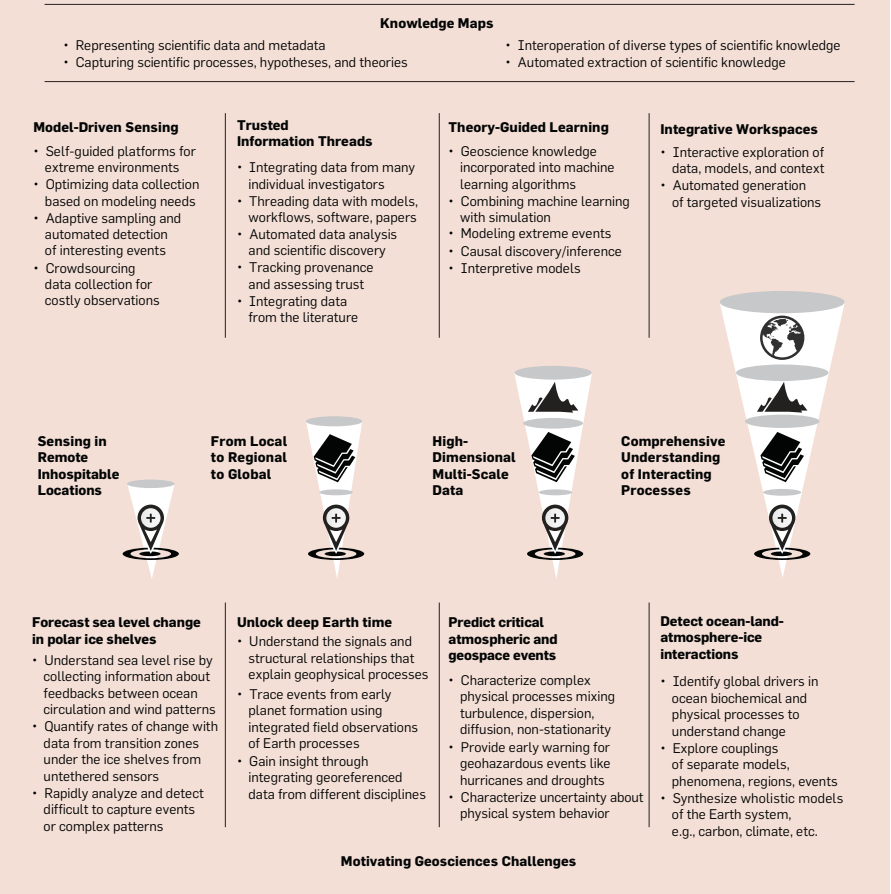
Earth systems phenomena are characterized by nonlinear, multiresolution, multi-scale, heterogeneous, and highly dynamic processes. Geosciences research is also challenged by extreme events and long-term shifts in Earth systems. The data available is intermittent, has significant sources of uncertainty, and is very sparse given the complexity and rich phenomena under study. Therefore, the small sample size of the datasets must be supplemented with the scientific principles underlying geosciences processes in order to guide knowledge discovery. For example, encapsulating knowledge about the physical processes governing Earth system datasets can help constrain the learning of complex nonlinear relationships in geoscience applications, ensuring theoretically consistent results. We need approaches that leverage the advances in data-driven research with methods that exploit the domain knowledge and scientific principles that govern the phenomena under study. These geoscience-aware systems will need to incorporate extensive knowledge about phenomena that combine physical, geological, chemical, biological, ecological, and anthropomorphic factors.

This body of research will lead to a new generation of *knowledge-rich intelligent systems* that contain rich knowledge and context in addition to data, enabling fundamentally new forms of reasoning, autonomy, learning, and interaction. The research challenges for creating knowledge-rich intelligent systems center on five major areas:

AI research.

New research in artificial intelligence (top) will result in a new generation of knowledge-rich intelligent systems that could address the significant challenges faced by geosciences (bottom). Knowledge-rich intelligent systems will exploit knowledge maps containing models and pre-existing knowledge in order to drive sensor data collection, create trusted information threads, power theory-guided learning, and enable integrative analytics.

A Research Agenda for Knowledge-Rich Intelligent Systems



1. *Knowledge representation and capture:* Capturing scientific knowledge about processes, models, and hypotheses.

2. *Sensing and robotics:* Prioritizing data collection based on the scientific knowledge available.

3. *Information integration:* Representing data and models as a “system of systems” where all knowledge is interconnected.

4. *Machine learning:* Enriching algorithms with knowledge and models of the relevant underlying processes.

5. *Interfaces and interactive systems:* Exploring and understanding user context using interconnected knowledge.

We describe these five areas in turn. For each area, we introduce major research directions followed by an overarching vision for that area.

Knowledge representation and capture. In order to create knowledge-rich

intelligent systems, scientific knowledge relevant to geoscience processes must be explicitly represented, captured, and shared.

Research directions:

1. *Representing scientific data and metadata.* Geoscientists are collecting more data than ever before, but raw data sitting on isolated servers is of little utility. Recent work on semantic and Linked Open Data standards enables publishing datasets in Web standard formats with open access licenses, creating links among datasets to further interoperability.² This leads to Web-embedded semantic networks and knowledge graphs that provide vast amounts of open interconnected knowledge about geosciences. Semantics, ontological representations, scientifically accurate concept mappings across domains, knowledge graphs,

and the application of Linked Open Data are all areas of active research to facilitate search and integration of data without a great deal of manual effort.⁵

2. *Capturing scientific processes, hypo-theses, and theories.* To complement the ontologies and data representations just discussed, a great challenge is representing the ever-evolving, uncertain, complex, and dynamic scientific knowledge and information. Important challenges will arise in representing dynamic processes, uncertainty, theories and models, hypotheses and claims, and many other aspects of a constantly growing scientific knowledge base. These representations need to be expressive enough to capture complex scientific knowledge, but they also need to support scalable reasoning that integrates disparate knowledge at different scales. In addition, scientists will need to understand the representations and trust the outcomes.

3. *Interoperation of diverse scientific knowledge.* Scientific knowledge comes in many forms that use different tacit and explicit representations: hypotheses, models, theories, equations, assumptions, data characterizations, and others. These representations are all interrelated, and it should be possible to translate knowledge fluidly as needed from one representation to another. A major research challenge is the seamless interoperation of alternative representations of scientific knowledge, from descriptive to taxonomic to mathematical, from facts to interpretation and alternative hypotheses, from smaller to larger scales, and from isolated processes to complex integrated phenomena.

4. *Authoring scientific knowledge collaboratively.* Formal knowledge representation languages, especially if they are expressive and complex, are not easily accessible to scientists for encoding understanding. A major challenge will be creating authoring tools that enable scientists to create, interlink, reuse, and disseminate knowledge. Scientific knowledge needs to be updated continuously, allow for alternative models, and separate facts from interpretation and hypotheses. These are new challenges for knowledge capture and authoring research. Finally, scientific knowledge should be created

collaboratively, allowing different contributors to weigh in based on their diverse expertise and perspectives.

5. *Automated extraction of scientific knowledge.* Not all scientific knowledge needs to be authored manually. Much of the data known to geoscientists is stored in semi-structured formats, such as spreadsheets or text, and is inaccessible to structured search mechanisms. Automated techniques are needed to identify and import these kinds of data into structured knowledge bases.

Research vision: Knowledge maps. We envision rich knowledge graphs that will contain explicit interconnected representations of scientific knowledge linked to time and space to form multidimensional *knowledge maps*. Interpretations and assumptions will be well documented and linked to observational data and models. Today's semantic networks and knowledge graphs link together distributed facts on the Web, but they contain simple facts that lack the depth and grounding needed for scientific research. Knowledge maps will have deeper spatiotemporal representations of processes, hypotheses, and theories and will be grounded in the physical world, interconnecting the myriad models of geoscience systems.

Robotics and sensing. Knowledge-informed sensing and data collection has great potential to do more cost-effective data gathering across the geosciences.

Research directions:

1. *Optimizing data collection.* Geoscience data is needed across many scales, both spatial and temporal. Since it is not possible to monitor every measurement at all scales all of the time, there is a crucial need for intelligent methods for sensing. New research is needed to estimate the cost of data collection prior to sensor deployment, whether that means storage size, energy expenditure, or monetary cost. A related research challenge is trade-off analysis of the cost of data collection versus the utility of the data to be collected.

2. *Active sampling.* Geoscience knowledge can be exploited to inform autonomous sensing systems to not only enable long-term data collection, but to also increase the effectiveness of sensing through adaptive sam-

pling, resulting in richer datasets at lower costs. Interpreting sensor data onboard allows autonomous vehicles to make decisions guided by real-time variations in data, or to react to unexpected deviations from the current physical model.

3. *Crowdsourcing data collection for costly observations.* Citizen scientists can contribute useful data (for example, collected through geolocated mobile devices) that would otherwise be very costly to acquire. One challenge in data collection through crowdsourcing is in ensuring high quality of data required by geoscience research. A potential area of research is to improve methods of evaluating crowdsourced data collection empirically, and to gain an understanding of the biases involved in the collection process.

Research vision: Model-driven sensing. New research on sensors will create a new generation of devices that will contain more knowledge of the scientific context for the data being collected. These devices will use that knowledge to optimize their performance and improve their effectiveness. This will result in new *model-driven sensors* that will have more autonomy and exploratory capabilities.

Information integration. Data, models, information, and knowledge are scattered across different communities and disciplines, causing great limitations to current geosciences research. Their integration presents major research challenges that will require the use of scientific knowledge for information integration.

Research directions:


1. *Integrating data from distributed repositories.* The geosciences have phenomenal data integration challenges. Most of the hard geoscience problems require that scientists work across sub-disciplinary boundaries and share very large amounts of data. Another facet of this issue is that the data spans a wide variety of modalities and greatly varying temporal and spatial scales. Distributed data discovery tools, metadata translators, and more descriptive standards are emerging in this context. Open issues include cross-domain concept mapping, entity resolution and scientifically valid data linking, and effective tools for finding, integrating, and reusing data.

2. *Threading scientific information and resources.* Scientific information and digital resources (data, software, models, workflows, papers, and so on) should be interconnected and interrelated according to their authors and use. Research challenges include developing new knowledge networks that accurately and usefully link together people, data, models, and workflows. This research will deepen our understanding of Earth science information interoperability and composition, and of how collaborative expertise and shared conceptual models develop.


3. *Automated data analysis and scientific discovery.* Capturing complex integrative data analysis processes as workflows facilitates reuse, scalable execution, and reproducibility. The pace of research could be significantly accelerated with intelligent workflow systems that automatically select data from separate repositories and carry out integrated analyses of data from different experiments. Through workflows that integrate large amounts of diverse data and interdisciplinary models, intelligent systems will lead to new discoveries.

4. *Tracking provenance and assessing trust.* Incoming data to the integration process must be analyzed for its fit and trustworthiness. The original sources must be documented, as well as the integration processes in order for the information to be understood and trusted. The challenges are in developing appropriate models and automating provenance/metadata generation throughout the integration and scientific discovery processes.

5. *Integrating data from the published literature.* Important historical data in geosciences is often only available in the published literature, requiring significant effort to integrate with new data. Text mining and natural language processing tools can already extract scientific evidence from articles.⁵ Important research challenges in this area include improving the quality of existing information extraction systems, minimizing the effort required to set up and train these systems, and making them scalable through the vast amounts of the published record. Another area of research is georeferencing extracted facts and integrating



This body of research will lead to a new generation of knowledge-rich intelligent systems that contain rich-knowledge and context in addition to data, enabling fundamentally new forms of reasoning, autonomy, learning, and interaction.



newly extracted information with existing data repositories.

Research vision: Trusted information threads. The proposed research will result in a scientifically accurate, useful, and trusted knowledge-rich landscape of data, models, and information that will include integrated broad-scale by-products derived from raw measurements. These products will be described to explain the derivations and assumptions to increase understanding and trust of other scientists. These *trusted information threads* will be easily navigated, queried, and visualized.

Machine learning. In order to address the challenges of analyzing sparse geosciences data given the complexity of the phenomena under study, new machine learning approaches that incorporate scientific knowledge will be needed so that inferences will be obtained better than from data alone.

Research directions:


1. *Incorporation of geoscience knowledge into machine learning algorithms.* Geoscience processes are very complex and high dimensional, and the sample size of the data is typically small given the space of possible observations. For those reasons, current machine learning methods are not very effective for many geoscience problems. A promising approach is to supplement the data with knowledge of the dominant geoscience processes.³ Examples from current work include the use of graphical models, the incorporation of priors, and the application of regularizers. Novel research is needed to develop new machine learning approaches that incorporate knowledge about geoscience processes and use it effectively to supplement the small sample size of the data. Prior knowledge reduces model complexity and makes it possible to learn from smaller amounts of data. Incorporating geoscience process knowledge can also address the high dimensionality that is typical of geoscience data. Prior knowledge constrains the possible relationships among the variables, reducing the complexity of the learning task.

2. *Combining machine learning and simulation approaches.* Machine learning offers data-driven methods to derive models from observational data. In contrast, geoscientists often use simulation models that are built.


Process-based simulation approaches impose conservation principals such as conservations of mass, energy, and momentum. Each approach has different advantages. Data-driven models are generally easier to develop. Process-based simulation models arguably provide reasonable prediction results for situations not represented in the model calibration period, while data-driven models are thought to be unable to extrapolate as well. Yet difficulties in the development of process-based simulation models, such as parameterization and the paucity of clear test results, can draw this claim into question. Intelligent Systems hold the promise of producing the evaluations needed to make the complex approaches used in data-driven and process-model simulation approaches more transparent and refutable. Such efforts will help to use these methods more effectively and efficiently. Novel approaches are needed that combine the advantages of machine learning and simulation models.

3. *Modeling of extreme values.* There are important problems in geosciences that are concerned with extreme events, such as understanding changes in the frequency and spatial distribution of extremely high temperature or extremely low precipitation in response to increase in greenhouse gas emissions. However, existing climate simulation models are often unable to reproduce realistic extreme values and therefore the results are not reliable. Although data science models offer an alternative approach, the heavy-tail property of the extreme values and its spatiotemporal nature poses important challenges to machine learning algorithms. A major challenge is presented by the spatiotemporal nature of the data.

4. *Evaluation methodologies.* Machine learning evaluation methodology relies heavily on gold standards and benchmark datasets with ground-truth labels. In geosciences there are no gold standard datasets for many problems, and in those cases it is unclear how to demonstrate the value of machine learning models. One possible approach involves making predictions, collecting observations, and then adjusting the models to account for differences between prediction and observations. Holding data mining competitions using such data would be a



Novel research is needed to develop new machine learning approaches that incorporate knowledge about geoscience processes and use it effectively to supplement the small sample size of the data.



very effective attractor for the machine learning community. Another alternative could be the creation of training datasets from simulations. Training datasets could be generated that would mimic real data but also have ground truth available, providing opportunity to rigorously train, test and evaluate machine learning algorithms.

5. *Causal discovery and inference for large-scale applications.* Many geoscience problems involve fundamental questions around causal inference. For example, what are the causes of more frequent occurrences of heat waves? What could be the causes for the change of ocean salinity? While it may be very hard to prove causal connections, it is possible to generate new (likely) hypotheses for causal connections that can be tested by a domain expert using methods such as generalization analysis of causal inference, causal inference in presence of hidden components, domain adaption and subsample data, Granger graphical models and causal discovery with probabilistic graphical models. Given the large amount of data available, we are in a unique position to use these advances to answer fundamental questions around causal inference in the geosciences.

6. *Novel machine learning methods motivated by geosciences problems.* A wide range of advanced machine learning methods could be effectively applied to geoscience problems. Moreover, geosciences problems drive researchers to develop entirely new machine learning algorithms. For example, attempts to build a machine learning model to predict forest fires in the tropics using multispectral data from earth observing satellites led to a novel methodology for building predictive models for rare phenomena¹ that can be applied in any setting where it is not possible to get high-quality labeled data even for a small set of samples, but poor-quality labels (perhaps in the form of heuristics) are available for all samples. Machine learning methods have already shown great potential in a few specific geoscience applications, but significant research challenges remain in order for those methods to be widely and easily applicable for other areas of geoscience.

7. *Active learning, adaptive sampling, and adaptive observations.* Many geoscience applications involve learn-

ing highly complex nonlinear models from data, which usually requires large amounts of labeled data. However, in most cases, obtaining labels can be extremely costly and demand significant effort from domain experts, costly experiments, or long time periods. Therefore, a significant research challenge is to effectively utilize a limited labeling effort for better prediction models. In machine learning, this area of research is known as active learning. Many relevant active sampling algorithms, such as clustering-based active learning, have been developed. New challenges emerge when existing active learning algorithms are applied in geosciences, due to issues such as high dimensionality, extreme events, and missing data. In addition, in some cases, we may have abundant labeled data for some sites while being interested in building models for other locations (for example, remote areas). Transfer active learning aims to solve the problem with algorithms that can significantly reduce the number of labeling requests and build an effective model by transferring the knowledge from areas with large amount of labeled data. Transfer active learning is still in the early stages and many opportunities exist for novel machine learning research.

8. *Interpretive models.* In the past few decades, we have witnessed many successes of powerful but complex machine learning algorithms, exemplified by the recent peak of deep learning models. They are usually treated as a black box in practical applications, but have been accepted by more communities given the rise of big data and their modeling power. However, in applications such as geosciences, we are interested in both predictive modeling and scientific understanding, which requires explanatory and interpretive modeling. A significant research area for machine learning is the incorporation of domain knowledge and causal inference to enable the design of interpretive machine learning approaches that can be understood by scientists and related to existing geosciences theories and models.

Research vision: Theory-guided learning. Geosciences data presents new challenges to machine learning approaches due to the small sample sizes relative to the complexity and non-linearity of

the phenomena under study, the lack of ground truth, and the high degree of noise and uncertainty. New approaches for *theory-guided learning* will need to be developed, where knowledge about underlying geosciences processes will guide the machine learning algorithms in modeling complex phenomena.

Intelligent user interaction. Scientific research requires well-integrated user interfaces where data can easily flow from one to another, and that include and exploit the user's context to guide the interaction. New forms of interaction, including virtual reality and haptic interfaces, should be explored to facilitate understanding and synthesis.

Research directions:

1. *Knowledge-rich context-aware recommender systems.* Scientists would benefit from proactive systems that understand the task at hand and make recommendations for potential next steps, suggest datasets and analytical methods, and generate perceptually effective visualizations. A major research challenge is to design recommender systems that appropriately take into account the complex science context of a geoscientist's investigation.

2. *Embedding visualizations throughout the science process.* Pervasive use of visualizations and direct manipulation interfaces throughout the science process would need to link data to hypotheses and allow scientists to experience models from completely new perspectives. These visualization-based interactive systems require research on the design and validation of novel visual representations that effectively integrate diverse data in 2D, 3D, multidimensional, multiscale, and multispectral views, as well as how to link models to the relevant data used to derive them.

3. *Intelligent design of rich interactive visualizations.* In order to be more ubiquitous throughout the research process, visualizations must be automatically generated and be interactive. One research challenge is to design visualizations. Another challenge is the design of visualizations that fit a scientist's problem. An important area of future research is the interactive visualizations and direct manipulation interfaces would enable scientists to explore data and

gain a better understanding of the underlying phenomena.

4. *Immersive visualizations and virtual reality.* There are new opportunities for low-cost usable immersive visualizations and physical interaction techniques that virtually put geoscientists into the physical space under investigation, while also providing access to other related forms of data. This research agenda requires bridging prior distinctions in scientific visualization, information visualization, and immersive virtual environments.

5. *Interactive model building and refinement through visualizations that combine models and data.* Interactive environments for model building and refinement would enable scientists to gain improved understanding on how models are affected by changes in initial data and assumptions, how model changes affect results, and how data availability affects model calibration. Developing such interactive modeling environments requires visualizations that integrate data with models, ensembles of models, model parameters, model results, and hypothesis specifications. These integrated environments would be particularly useful for developing machine learning approaches to geosciences problems, for example in assisting with parameter tuning and selecting training data. A major challenge is the heterogeneity and complexity of these different kinds of information that needs to be represented.

6. *Interfaces for spatiotemporal information.* The vast majority of geosciences research products is geospatially localized and with temporal references. Geospatial information requires specialized interfaces and data management approaches. New research is needed in intelligent interfaces for spatiotemporal information that exploit the user's context and goals to identify implicit location, to disambiguate textual location specification, or to decide what subset of information to present. The small form factor of mobile devices is also constraint in developing applications that involve spatial data.

7. *Collaboration and assistance for data analysis and scientific discovery processes.* Intelligent workflow systems could help scientists by automating routine aspects of their work.

Because each scientist has a unique workflow of activities, and because their workflow changes over time, a research challenge is that these systems need to be highly flexible and customizable. Another research challenge is to support a range of workflows and processes, from common ones that can be reused to those that are highly exploratory in nature. Such workflows systems must enable collaborative design and analysis and be able to coordinate the work of teams of scientists. Finally, workflow systems must also support emerging science processes, including crowdsourcing for problems such as data collection and labeling.

Research vision: Integrative workspaces. New research is required to allow scientists to interact with all forms of knowledge relevant to the phenomenon at hand, to understand uncertainties and assumptions, and to provide many alternative views of integrated information. This will result in user interfaces focused on *integrative workspaces*, where visualizations and manipulations will be embedded throughout the analytic process. These new intelligent user interfaces and interaction modalities will support the exploration not only of data but of the relevant models and knowledge that provide context to the data. Research activities will flow seamlessly from one user interface to another, each appropriate to the task at hand and rich in user context.

Conclusion


This article presented research opportunities in knowledge-rich intelligent systems inspired by geosciences challenges. Crucial capabilities are needed that require major research in knowledge representation, selective sensing, information integration, machine learning, and interactive analytics.

Enabling these advances requires intelligent systems and geosciences researchers work together to formulate knowledge-rich frameworks, algorithms, and user interfaces. Recognizing that these interactions are not likely to occur without significant facilitation, a new Research Coordination Network on Intelligent Systems for Geosciences has been created to enable sustained communication across

these fields that do not typically cross paths. This network focuses on three major goals. First, the organization of joint workshops and other forums will foster synergistic discussions and collaborative projects. Second, repositories of challenge problems and datasets with crisp problem statements will lower the barriers to getting involved. Third, a curated repository of learning materials to educate researchers and students alike will reduce the steep learning curve involved in understanding advanced topics in the other discipline. Additionally, members of the Research Coordination Network are engaging other synergistic efforts, programs, and communities, such as artificial intelligence for sustainability, climate informatics, science gateways, and the U.S. NSF Big Data Hubs.

A strong research community in this area has the potential to have transformative impact in artificial intelligence research with significant concomitant advances in geosciences as well as in other science disciplines, accelerating discoveries and innovating how science is done.

Acknowledgments

This work was sponsored in part by the Directorate for Computer and Information Science and Engineering (CISE) and the Directorate for Geosciences (GEO) of the U.S. National Science Foundation under awards IIS-1533930 and ICER-1632211. We thank NSF CISE and GEO program directors for their guidance and suggestions, in particular Hector Muñoz-Avila and Eva Zanzerkia for their guidance, and Todd Leen, Frank Olken, Sylvia Spengler, Amy Walton, and Maria Zemankova for suggestions and feedback. We also thank all the participants in the Research Coordination Network on Intelligent Systems for Geosciences for creating the intellectual space for productive discussions across these disciplines. 

References

1. Gil, Y. and Pierce, S. (Eds). Final Report of the 2015 NSF Workshop on Information and Intelligent Systems for Geosciences. National Science Foundation Workshop Report, October 2015; <http://dl.acm.org/collection.cfm?id=C13> and <http://is-geo.org/>
2. Berners-Lee, T. Linked data. *Design Issues* (retrieved Nov. 11, 2017); <https://www.w3.org/DesignIssues/LinkedData.html>
3. Karpatne, A. et al. Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Transactions on Knowledge and Data Engineering* 29, 10 (2017) 2318–2331.

4. Mithal, V., Nayak, G., Khandelwal, A., Kumar, V., Oza, N.C. and Nemari, R. RAPT: Rare class prediction in absence of true labels. *IEEE Transactions on Knowledge and Data Engineering*, 2017; DOI: 10.1109/TKDE.2017.2739739.
5. Narock, T. and Fox, P. The Semantic Web in Earth and space science. Current status and future directions. *Studies in the Semantic Web*. IOS Press, 2015.
6. National Research Council, Committee on Challenges and Opportunities in the Hydrologic Sciences, Water Science and Technology Board, Division on Earth and Life Studies. Challenges and Opportunities in the Hydrologic Sciences. National Academies Press, Washington, D.C., 2012, 188. ISBN: 978-0-309-22283-9.
7. National Research Council, Committee on a Decadal Strategy for Solar and Space Physics (Heliophysics); Space Studies Board; Aeronautics and Space Engineering Board; Division of Earth and Physical Sciences. Solar and Space Physics: A Science for a Technological Society. National Academies Press, Washington, D.C., 2013, 466. ISBN 978-0-309-16428-3.
8. National Research Council, Committee on Guidance for NSF on National Ocean Science Research Priorities: Decadal Survey of Ocean Sciences, Ocean Studies Board; Division on Earth and Life Studies. Sea Change: 2015-2025 Decadal Survey of Ocean Sciences. National Academies Press, Washington, D.C., 2014, 98. ISBN 978-0-309-36888-5.
9. National Research Council, Committee on New Research Opportunities in the Earth Sciences. New Research Opportunities in the Earth Sciences at the National Science Foundation. National Academies Press, Washington, D.C., 2012, 216. ISBN 978-0-309-21924-2.
10. National Research Council, Committee to Review the NSF AGS Science Goals and Objectives. Review of the National Science Foundation's Division on Atmospheric and Geospace Sciences Goals and Objectives Document. National Academies Press, Washington, D.C., 2014, 36. ISBN 978-0-309-31048-2.
11. National Science Foundation. Dynamic Earth: GEO Imperatives and Frontiers 2015–2020. Advisory Committee for Geosciences, 2014.
12. Peters, S.E., Zhang, C., Livny, M. and Ré, C. A machine reading system for assembling synthetic paleontological databases. *PLoS ONE* 9, 12 (2014).

Yolanda Gil, University of Southern California; **Suzanne A. Pierce**, The University of Texas Austin; **Hassan Babaie**, Georgia State University; **Arindam Banerjee**, University of Minnesota; **Kirk Borne**, Booz Allen Hamilton; **Gary Bust**, Johns Hopkins University; **Michelle Cheatham**, Wright State University; **Imme Ebert-Uphoff**, Colorado State University; **Carla Gomes**, Cornell University; **Mary Hill**, University of Kansas; **John Horel**, University of Utah; **Leslie Hsu**, Columbia University; **Jim Kinter**, George Mason University; **Craig Knoblock**, University of Southern California; **David Krum**, University of Southern California; **Vipin Kumar**, University of Minnesota; **Pierre Lermusiaux**, Massachusetts Institute of Technology; **Yan Liu**, University of Southern California; **Chris North**, Virginia Tech; **Victor Pankratius**, Massachusetts Institute of Technology; **Shanan Peters**, University of Wisconsin-Madison; **Beth Plale**, Indiana University Bloomington; **Allen Pope**, University of Colorado Boulder; **Sai Ravela**, Massachusetts Institute of Technology; **Juan Restrepo**, Oregon State University; **Aaron Ridley**, University of Michigan; **Hanan Samet**, University of Maryland; **Shashi Shekhar**, University of Minnesota; **Katie Skinner**, University of Michigan; **Padhraic Smyth**, University of California Irvine; **Basil Tikoff**, University of Wisconsin-Madison; **Lynn Yarmey**, National Snow and Ice Data Center; **Jia Zhang**, Carnegie Mellon University

Correspondence regarding this article should be directed to **Yolanda Gil** (gil@isi.edu).

Copyright held by authors/owners.



Watch the authors discuss this work in the exclusive *Communications* video. <https://cacm.acm.org/videos/intelligent-systems-for-geosciences>