*CCC Whitepaper*

**July 2025**

# The Imperative for Grand Challenges in Computing

**Authors**

David Jensen (University of Massachusetts, Amherst), Daniel Lopresti (Lehigh University), Mary Lou Maher (Computing Research Association), Manish Parasar (University of Utah), Rajmohan Rajaraman (Northeastern University), William Regli (University of Maryland), Mona Singh (Princeton University), Holly Yanco (University of Massachusetts, Lowell)

## Introduction

Computing is now an indispensable component of nearly all technologies and is ubiquitous for vast segments of society. It is also essential to discoveries and innovations in most disciplines. However, while past grand challenges in science have involved computing as one of the tools to address the challenge, they have not been principally about computing. In this paper, we ask: What are the grand challenges in computing, and how would we go about identifying such challenges? In other words, what are the grand challenges of computing that would inspire transformative research with the potential to dramatically improve the future of our society?

Climate, health, energy, national security, and the acceleration of scientific inquiry and economic well-being are some areas that provide context for challenge problems. These multi-disciplinary problems necessitate the creation of use-inspired and translational research in search of solutions.

We see a significant benefit in taking a more intentional approach as a field. While the pace of development in computing over the past few decades has been exhilarating, too often the implications of our progress do not become apparent until after large investments have been committed and the impacts on science and the public are not clear or as intended. We believe that taking a long view that is intentional about translation and realizing impacts is appropriate as our field continues to mature.

## Why Think About Grand Challenges and Why Now?

While change may seem to appear suddenly, computing revolutions often emerge over the course of several generations of foundational, use-inspired, and translational research activity, perhaps unfolding over many years until critical tipping points are reached (representative examples include the Internet, human-computer interfaces, deep learning, large language models, various cybersecurity solutions, open source software). While there may be some key actors, the conditions for success are usually the product of many individuals and organizations contributing multiple, varying disciplinary perspectives. The development of the ARPAnet and Internet through the 1970s, and the emergence of the World Wide Web in the late 1990s are illustrations of such collective innovations. Compelling applied and use-inspired objectives create the need to bring these emerging pieces together. Examples range from the development of IBM 360 for general purpose business computing to the "Mother of All Demos" of Douglas Engebart in 1968 showing integrated collaboration and software development, and the iPhone in 2007 highlights use-inspired objectives. Often, computing revolutions come from major *intentional* and *translational* investments in infrastructure, testbeds, and challenge problems to motivate and focus the community. Examples include DARPA's Strategic Computing Initiative of the 1980s and National Science Foundation's "10 Big Ideas," which often included shared resources and testbeds. The supercomputer and cyberinfrastructure facilities managed by the Department of Energy and the NSF serve a similar purpose.

It is important and timely to think strategically about these trends and the role of computer science research. It has been noted often, and by others, that the invention of calculus and the establishment of the scientific method in the 17th century created the framework for the industrial revolution and the past three centuries of progress enabled by the field of engineering. This century is the century of computing machines, and computer science, which is now 70 years old, and has firmly established a new framework that can enable us to reason about and address the challenges and opportunities facing society. With the rapid and wide-ranging advances in artificial intelligence, we believe computer science is at an inflection point motivating us to ask: Given the state of the world today, and of science and technology, what are the grand challenges uniquely centered on "computing research" that can catalyze the next generation of discoveries?

CCC
Computing Community Consortium
Catalyst

## A Recent History of Grand Challenges

The concept of “grand challenges” is not new and over the last centuries several have changed the course of history. In 1714, the British government established the Longitude Act[1], offering a prize for a means of accurately calculating the longitude of ships at sea. This prize was won by amateur horologist John Harrison and his H4 timepiece—forerunner of our Global Position System today. The war of 1812 made the US Army aware of the benefits of interchangeable parts, leading to several decades of innovation in measurement and manufacturing. In 1900, David Hilbert’s stated 27 fundamental mathematical problems[2], the topics of which became the focus of mathematics for over half of the 20th century. New York City hotelier Raymond Orteig[3] established a prize in 1919 for the first powered flight over the Atlantic Ocean—eventually won in 1927 by Charles Lindbergh. World War II precipitated several technological challenges that led to innovations such as the proximity fuse[4], radar[5], and the atomic bomb. The Cold War gave rise to the space race and the Apollo Program. The concepts of “Manhattan Project” and “Moonshot” are now synonymous with our notion of a “Grand Challenge”.

There have been grand challenges in a variety of other areas of science and engineering. Examples of such “grand challenges” might include:

- Design and launch of various space telescopes (Hubble Space Telescope (HST), 1970s–present; Kepler Space Telescope, 2000–2018; Spitzer Space Telescope, 1990–2020; James Webb Space Telescope (JWST), 1996–present).
- The Human Genome Project (planned in 1984, started in 1990, and fully completed in January 2022) created a complete map of human DNA.
- Search for the Higgs Boson and development of the Large Hadron Collider (1998-2010).
- The development of gravitational astronomy and the construction of the Laser Interferometer Gravitational-Wave Observatory (LIGO)[6].
- The Clay Mathematics Foundation maintains a list of Hilbert-like Millennium Prize challenges[7] in various areas of mathematics.

---

[1] https://en.wikipedia.org/wiki/Longitude_Act
[2] https://en.wikipedia.org/wiki/Hilbert%27s_problems
[3] https://en.wikipedia.org/wiki/Orteig_Prize
[4] https://en.wikipedia.org/wiki/Proximity_fuze
[5] https://en.wikipedia.org/wiki/History_of_radar
[6] https://en.wikipedia.org/wiki/LIGO
[7] https://en.wikipedia.org/wiki/Millennium_Prize_Problems

Perhaps most related, the National Academy of Engineering organized a multi-year planning effort that culminated in 2013 with “14 Grand Challenges for Engineering for the 21st Century”[8], which included topics such as “Manage the Nitrogen Cycle” and “Prevent Nuclear Terror”. The federal government has often taken a leadership role in this process; recent examples include the 21st Century Grand Challenges[9], the Cancer Moonshot in 2016, and the National Quantum Initiative Act of 2018.

Computing, in various forms, is an essential element of these “grand challenges” (especially the recent ones). But, while computing technologies are essential in these areas, the challenges are, aside from “P vs NP” in the Millennium Prize list, not principally about computing itself. What can we learn from past "grand challenges" efforts, specifically within the field of computing?

Initiated in the 1980s, DARPA’s Strategic Computing Initiative (SCI)[10] outlined the framework for the computer science innovations and advancements that emerged in the subsequent decades. Initial “challenges” were related to autonomous driving, language understanding, and human-machine teaming (i.e., “the pilot’s assistant”). These came to fruition in a variety of ways. DARPA’s Grand Challenges in 2004, 2007, and 2012-2015 pushed the frontiers on fully autonomous ground vehicles and autonomous emergency-maintenance robots—technologies we now see in commercial practice. Siri[11] emerged from DARPA programs (most specifically the Cognitive Assistant that Learns and Organizes (CALO) program) in the early 2000s as did innovations in human language. The leadership mantles for these areas were assumed by industry, leading to numerous commercial innovations we now find ubiquitous.

In 2002, Computing Research Association hosted a significant “grand challenges” conference[12] that, in addition to many computing luminaries among the attendees, included several of the literal founders of the field of computer science. Contemporaneously, there were similar efforts in the United Kingdom to identify computing “grand challenges”. This 2002 conference produced five candidate challenges: (1) Systems you can count on; (2) A teacher for every learner; (3) 911.net (ubiquitous information systems); (4) Augmented cognition; and (5) Conquering complexity. Assessing progress since that time is difficult due to, for the most part, the lack of specificity and metrics associated with these challenges. In at least two of the

---

[8] https://www.engineeringchallenges.org/challenges.aspx
[9] https://obamawhitehouse.archives.gov/administration/eop/ostp/grand-challenges
[10] https://en.wikipedia.org/wiki/Strategic_Computing_Initiative and https://mitpress.mit.edu/9780262529266/strategic-computing/
[11] https://en.wikipedia.org/wiki/Siri
[12] https://archive.cra.org/Activities/grand.challenges/index.html

cases, the objectives are subjective, making it difficult to know exactly when one has been achieved.

In addition to SCI and CRA's efforts, there have been additional computational challenges sponsored by NSF, DARPA, and others, including those in focused topics such as cybersecurity, high-performance computing, and, of course, AI. What lessons might we learn from past efforts at "grand challenges"? We observe three things.

First, for some, there is a profound scientific question for which the challenge is to design and build *profoundly new kinds of scientific instruments* to understand and ultimately answer the question. This is seen most obviously in the various programs that have built telescopes and other observational capabilities (e.g., LIGO, JWSC,.etc.). In computing, the goal of creating an operational Quantum Computer might be seen as a similar example.

Second, most of these challenges (certainly including those that build scientific instruments) have *tangible and definite outcomes*. Like the Manhattan Project or the original Moonshot, challenges related to mathematical proof (Hilbert's questions and the Millennium Prize) certainly fall into this category, i.e., once there is a proof of "P vs NP", the challenge is achieved. The same can also be said for the Human Genome Project, Operation Warp Speed, and the development of the COVID-19 vaccines, and any other challenge with a specific engineering objective.

Lastly, an effective challenge, even one that lacks a tangible and definite goal, at least has *metrics* by which one can measure progress. Self driving vehicles have seen a steady increase in the complexity of measurement: from driving across a desert, to passing the California driver's test, to fully autonomous driving. The current quest for Artificial General Intelligence (AGI) might also be viewed in this way, measured as successive improvements against various benchmarks—indeed, the very disagreement around measurements and metrics for AI systems may indeed be a fundamental part of the challenge itself.

## How to Think About Grand Challenges in Computing?

As we explore grand challenges in computing, it is essential to develop a framework for conceptualizing and advancing such challenges. Such a framework must enable holistic thinking about computing research from different perspectives; in addition to *scientific and societal* grand challenges where *computing research is essential*, the framework should enable the formulation of grand challenges for computer science as a field, including the acceleration and amplification of its translational impacts. Toward this end, we posit that every grand challenge must satisfy the following key criteria.

- **Impact:** The challenge should address a critical need or urgent societal problem.
- **Ambition:** The challenge should be ambitious and have the potential to drive significant progress in the computing field.
- **Feasibility:** The challenge should be beyond the reach of current computing techniques but perhaps barely feasible using disruptive new ideas on the horizon.
- **Interdisciplinarity:** The challenge should draw expertise from multiple computing research areas and stakeholders from diverse backgrounds.
- **Measurability:** Progress towards addressing the challenge should be measurable. While a formal quantification of progress may not always be possible at the stage of formulating the challenge, it is critical to articulate outcomes that will result from a successful completion of the challenge.

It is important that the process for formulating grand challenges brings together a broad spectrum of researchers, practitioners, and stakeholders, from within computing and beyond. Such an approach will create space for brainstorming on the challenges within the framework of the above criteria, and can identify grand challenge prototypes which may be expressed by brief answers to the following questions:

*What is the challenge?*

*Why now and what makes this challenge "grand" and (barely) feasible?*

*How will we know if we succeeded?*

## Computing Grand Challenge Prototypes: A Sampling

The Grand Challenges Task Force of the CCC Council organized four virtual roundtable sessions in Spring 2024 with 36 attendees, which included researchers across the computer science spectrum. In July 2024, the Task Force led a dynamic 90-minute session at the CRA Summit in Snowbird with over 70 attendees, drawn from leadership in academia and industry research. More recently, in May 2025, the Task Force held a grand challenges session at the CIFellows Symposium with 40 early-career computing researchers.

In the following, we present a few proposals for grand challenge prototypes, which originated from these discussions. We emphasize that these proposals are preliminary. While our framework guided the formulation, the proposals do not fully meet the criteria laid out above. We need more thought and discussion within the research community before formalizing these as grand challenges.

## Data-Driven Personalized Healthcare

### *What Is the Challenge?*

Transform how doctors deliver healthcare by considering personal data (e.g., blood-work labs, fitness trackers, symptoms, existing diagnoses and medicines, genomic information) in the context of data collected for all individuals seeking healthcare across the country/world in order to provide personalized treatments and recommendations. A closely intertwined challenge is to develop robust and reliable digital twins of individual humans to enhance predictive accuracy in preventative medicine.

### *Why Now, and What Makes This Challenge "Grand" and (Barely) Feasible?*

The challenge requires solving a number of pressing and vexing problems in diverse areas of computing, including privacy and security (sharing and computing on sensitive data), programming languages (interoperability of information arising from distinct hospital systems/coding systems), AI (recognizing shared patterns in data, foundation models), and computational biology and medicine (genomic information, personalized medicine, pharmacogenomics). The time seems right for this challenge with the confluence of the scale and diversity of individual health data, continuing advances in evidence- and data-driven medicine, and increased deployment of AI tools in healthcare.

### *How Will We Know If We Succeeded?*

Precision health solutions will be universally adopted in the healthcare industry, and predictive errors will be extremely rare. Another important metric is improvement in individual health measures and outcomes.

## Ubiquitous Data Provenance

### *What Is the Challenge?*

Develop technology to facilitate reliable, accurate, and ubiquitous tracing of the sources for all public data, including scientific datasets, photographs, videos, articles, and other internet artifacts. This would be an important tool for identifying AI-generated content, assigning due credit, and tracking data sources and transformations.

### *Why Now, and What Makes This Challenge "Grand" and (Barely) Feasible?*

This interdisciplinary challenge needs new techniques in many computing areas, including generative AI, cryptography, databases, information retrieval, human-computer interaction, and systems. Specific questions concerning digital signatures of LLM systems are already being studied, but the challenges are immense and span both the socio-technical side as well as the foundational aspects of how data is stored, retrieved, and processed.

### *How Will We Know If We Succeeded?*

Provenance of public data will be accurate and complete, and the technology will be widely adopted.

## DIY Software

### *What Is the Challenge?*

Develop technology for designing effective and reliable software systems where the user does not need to be a programmer or have any knowledge of programming. Compute power can and should be leveraged by more people than just computer scientists. We want technology development to be available to people from all kinds of fields, like an English teacher, to create and use tools that are easy to understand. By doing this, we can open up innovation to more people and benefit from their diverse ideas and knowledge.

### *Why Now, and What Makes This Challenge "Grand" and (Barely) Feasible?*

Meeting this challenge requires new techniques in natural-language processing and AI that transcend domains, advances in HCI that cover users with widely varying technological skills, innovations in software design that introduce levels of abstraction allowing anyone to create a personalized system for a specific purpose. The ongoing advances in generative AI can provide a springboard to efforts toward this challenge.

### *How Will We Know If We Succeeded?*

Success will be measured by the extent of adoption of the technology, the reliability of the relevant tools, and the boost to productivity.

## Net Zero Computing

### *What Is the Challenge?*

Develop computing hardware, systems software, algorithms, and applications that optimize the tradeoffs between the costs of the resources they consume and the value of the tasks they are designed to accomplish. It will also be necessary to conceive the corresponding theoretical foundations that allow for the analysis and quantification of the resource usage so that these tradeoffs can be managed, similar to what asymptotic time/space complexity did for the computing field beginning in the 1970s. By "costs," we are including societal impacts in addition to business costs, and by "resources" we are not only referring to power usage, but all of the resources needed to design, manufacture, distribute, and operate computer hardware and software.

### *Why Now and What Makes This Challenge "Grand" and (Barely) Feasible?*

Recent technological developments have provided some promising approaches that could form part of the solution. For example, thanks to the cloud, computation is fully and transparently portable with no apparent added latency or costs, unlike many other production activities, including traditional manufacturing. Beyond what industry and governments can accomplish, the computing research community will play a key role because succeeding at this challenge requires adopting a different mindset toward the way we conduct and evaluate our research, as well as the way we educate our students.

### *How Will We Know If We Succeeded?*

Users and developers of computing will have full knowledge of the cost-benefit tradeoffs of their applications and, as a result, will be in a position to make better-informed choices between the myriad of options available to them.

## Personal Teacher / Coach

### *What Is the Challenge?*

Develop an AI tool with which anyone at all – including an individual of limited technical background – can create and customize a personal teacher/coach that will aid the user in learning any new skill of interest faster and more effectively than working with a traditional human teacher/coach. One domain would be applications that require some physical skill (e.g., learning to play the piano or to play tennis, or to paint landscapes in a certain artistic style), while another might be purely thought-based (e.g., learning probability and statistics, or how to write a poem or a Hollywood screenplay).

### *Why Now and What Makes This Challenge "Grand" and (Barely) Feasible?*

AI, including machine learning, computer vision, and large language models, has advanced to the point that machines now rival human performance when quantified using relatively simplistic measures such as raw accuracy. Substantial work, including pure research in a variety of areas, will be needed to integrate these isolated capabilities into useful complete systems that can guide human users as adeptly as a talented teacher or coach.

### *How Will We Know If We Succeeded?*

The learning of new skills will be completely transformed from traditional hit-or-miss methods designed for the "average" user to much more personalized approaches that permit everyone to achieve their full potential across a wide variety of topics that are both interesting and valuable to them.

## Make CS and AI Research Trustworthy and Explainable

### *What Is the Challenge?*

Develop the theory and systems for making research artifacts in computer science trustworthy and explainable by verifying the proofs of all theoretical claims, ensuring the reproducibility of experimental results, and validating statistical data analyses. The need is especially pressing given the widespread use of machine learning in scientific research and the explosive growth of generative AI tools.

### *Why Now, and What Makes This Challenge “Grand” and (Barely) Feasible?*

Solving this challenge requires new techniques in algorithms and proof assistants (for validating proofs), systems, programming languages, and software engineering (for reproducing experimental results), AI, NLP, and data science (for information extraction and statistical validity). These advances can be fueled by impressive recent progress in formal methods (including theorem-provers and SAT solvers), and LLMs capable of processing and generating multi-modal data. The impact to society is through the assurance gained by the correctness of computing research outputs. The impact can be more grand and more direct if the solutions are extended to other sciences.

### *How Will We Know If We Succeeded?*

Success is when the verification of CS research artifacts is widely adopted, and the publication of results and deployment of artifacts are accompanied by certificates that establish their validity.

## Conclusion

Advances in computing are essential for addressing the urgent challenges in science and society. It is thus important that we are intentional in where we focus and how we conduct research in computer science. As the past has shown us, grand challenges can play a critical role in defining research directions and catalyzing the research community. In this paper, we have motivated and developed a framework for identifying computing grand challenges underlying the most urgent societal problems. This is a first step in this process, and we call on the community to advance and refine this framework and build on it to formulate their visions for computing grand challenges.

---


## Suggested Citation

Jensen, D., Lopresti, D.,  Maher, M., Parashar, M., Rajaraman, R., Regli, W., Singh, M., Yanco, H., The Imperative for Grand Challenges in Computing. Washington, D.C.: Computing Research Association (CRA), July 2025.
[http://cra.org/ccc/wp-content/uploads/sites/2/2025/07/The-Imperative-for-Grand-Challenges-in-Computing.pdf](http://cra.org/ccc/wp-content/uploads/sites/2/2025/07/The-Imperative-for-Grand-Challenges-in-Computing.pdf)

# ACKNOWLEDGMENTS

**Reviewers**

CCC and the authors acknowledge the reviewer whose thoughtful comments improved the report:

- David Danks, University of California, San Diego

## U.S. National Science Foundation

The CCC AI Futures Task Force was supported by the Computing Community Consortium through the National Science Foundation under Grant No. 2300842. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

## About the Computing Community Consortium (CCC)

A programmatic committee of the Computing Research Association (CRA), CCC enables the pursuit of innovative, high-impact computing research that aligns with pressing national and global challenges. Of, by, and for the computing research community, CCC is a responsive, respected, and visionary organization that brings together thought leaders from industry, academia, and government to articulate and advance compelling research visions and communicate them to stakeholders, policymakers, the public, and the broad computing research community.

CCC
Computing Community Consortium
Catalyst