Report of the National Science Foundation Workshop on Information Theory and Computer Science Interface

October 17-18, 2003

Chicago, Illinois

1 Panelists

The following are the attendees (listed in alphabetical order) of the National Science Foundation (NSF) Workshop on Information Theory and Computer Science Interface held in Chicago, Illinois, during October 17 to 18, 2003.

Name	Affiliation
Toby Berger	Cornell University
Tom Cover	Stanford University
Kevin Compton	University of Michigan
Michelle Effros	California Institute of Technology
Tony Ephremides	University of Maryland
Bruce Hajek	University of Illinois, Urbana-Champaign
Philippe Jacquet	INRIA
John Kieffer	University of Minnesota
Michael Mitzenmacher	Harvard University
Alon Orlitsky	University of California, San Diego
Nicholas Pippenger	Princeton University
Serap Savari	University of Michigan
Gadiel Seroussi	Hewlett-Packard Laboratories
James Storer	Brandeis University
Madhu Sudan	Massachusetts Institute of Technology
Wojciech Szpankowski	Purdue University
Vahid Tarokh	Harvard University
Sergio Verdu	Princeton University
Jeff Vitter	Purdue University
Marcelo Weinberger	Hewlett-Packard Laboratories
En-hui Yang	University of Waterloo
Zhen Zhang	University of Southern California

Members of the organizing committee of the workshop were:

- John Kieffer, University of Minnesota
- Wojciech Szpankowski, Purdue University (Chair)
- En-hui Yang, University of Waterloo

2 Executive Summary

The NSF workshop on Information Theory and Computer Science Interface was held in Chicago, Illinois, during October 17 to 18, 2003. This is perhaps the first NSF sponsored workshop addressing fundamental research issues and research training in the interface of information theory (IT) and computer science (CS), even though the interplay between IT and CS dates back to the founding father of information theory, Claude E. Shannon. Twenty two (22) leading researchers and innovators from both academia and industry participated in the workshop to express their views and recommendations on

- fundamental research issues in the interplay of IT and CS,
- future directions and challenges of IT and CS,
- radical innovations in communications and computation taking place at the crossroad of IT and CS,
- frameworks for future cross-cutting research activities in IT and CS,
- training of highly qualified personnel in the interaction of IT and CS, including graduate and undergraduate student training, and their corresponding curriculum developments, and
- better use of NSF's limited resources to promote joint IT and CS excellence in the areas of research, research training, and research leadership.

Innovation opportunities, research challenges, frameworks of interdisciplinary collaboration, and training improvements were identified and agreed upon in key areas of the interaction between IT and CS. These include networking (wireless networks, sensor networks, ad hoc networks, etc.), coding, compression, Shannon theory, theoretic computer science, as well as emerging areas, such as biology, coding for nano-structure, and digital watermarking.

We recommend a new initiative by NSF on basic research that exploits the interface between Computer Science and Information Theory. The Information Technology revolution that has affected Society and the world so fundamentally over the last few decades is squarely based on computation and communication, the roots of which are respectively Computer Science and Information Theory. The current organization of the National Science Foundation makes it difficult to identify the programs that could harbor and support research efforts on the foundations and interfaces of these two core disciplines. So, a special initiative by the Foundation would rekindle research on the many maturing topics outlined in this report that point to the strong interplay between Computer Science and Information Theory.

3 Introduction and Motivation

The history of scientific discovery is a testimony to the success of multidisciplinary research that couples disparate branches of science. The recent proof of Fermat's last theorem by A. Wiles provides us with a prime example of this. Recall that Wiles' success followed on the heels of the works of Frey and Ribert that connected the seemingly unrelated Taniyama-Shimura conjecture (of modular forms) and Fermat's last theorem. Along the same vein, Feynman's fiddling with the axial wobble of a cafeteria plate at Cornell led him to his version of quantum mechanics. Feynman's "sum over histories" approach was picked up by mathematician Mark Kac, who wrapped it in the now-famous Feynman-Kac path integral. Likewise, a proof of the prime number conjecture had eluded mathematicians for a hundred years, until Hadamard and de la Vallee-Poussin applied complex analysis to prove it. In fact, Hadamard illustrated the utmost importance of complex analysis with his claim that the shortest path between two truths on the real line passes through the complex plane. Most recently, the assembly of the human genome by a group of biologists, computer scientists, mathematicians, and physicists is the ultimate evidence supporting multidisciplinary research.

The interplay between IT and CS dates back to the founding father of information theory, Claude E. Shannon. His landmark paper "A Mathematical Theory of Communication" is hailed as the foundation for information theory. Shannon also worked on problems in computer science such as chess-playing machines (or algorithms, in the currently popular terminology) and computability of different Turing machines.

Ever since Shannon's work on both information theory and computer science, the research in the interplay between IT and CS has continued and expanded in many exciting ways. In late 1960s and early 1970s, there were tremendous interdisciplinary research activities from IT and CS, exemplified by the work of Kolmogorov, Chaitin, and Solomonoff, with the aim of establishing the algorithmic information theory. Motivated by approaching the Kolmogorov complexity algorithmically, A. Lempel (a computer scientist), and J. Ziv (an information theorist) worked together in later 1970s to develop compression algorithms that are now widely referred to as Lempel-Ziv algorithms. Today, these are a de facto standard for lossless text compression; they are used pervasively in computers, modems, and communication networks.

Lempel-Ziv algorithms have been (and will certainly continue to be) constant inspirations for both IT and CS communities; researchers in both these fields are still working together towards a more fundamental understanding of data compression. For example, recent research in the theoretical computer science subfield (mostly from the Analysis of Algorithms group http://pauillac.inria.fr/algo/AofA/index.html) contributed to solutions of some open problems in data compression (e.g., the limiting distribution of the number of phrases in the LZ'78 scheme, precise bounds on redundancy of the LZ'78 scheme). In 1999, by using concepts from CS, Kieffer and Yang provided a simple technique to bound the pointwise redundancy of LZ'78. In 2004, a seminar and a graduate course on analytic methods in information theory and analysis of algorithms took place at MSRI, Berkeley.

The formal language approach developed over the last 50 years by computer scientists and logicians helped Kieffer and Yang to create a new universal lossless source coding theory called grammar-based coding theory. It provides a framework to design universal lossless data compression algorithms by integrating the power of pattern matching and context modeling. In the grammar-based coding theory, each algorithm first transforms a data sequence into a grammar, and then uses arithmetic coding to compress the grammar. Interestingly, using grammars for data compression and for extracting patterns from DNA was also explored independently by a group of computer scientists (including Craig Nevill-Manning, Ian Witten, Alberto Apostolico, and others) at about the same time. Pattern matching (e.g., a lossy extension of the Lempel-Ziv scheme) was also used for video compression. Grammar-based coding and other compression algorithms have also been used by Lanctot, Li, and Yang to estimate the entropy of DNA sequences and analyze genome sequences; note that both Lanctot and Li are computer scientists.

Witnessing a proliferation of research which has benefited from the interplay between IT and CS—each of these two fields has impacted the other in providing design paradigms and methods for obtaining performance bounds—*IEEE Transactions on Information Theory* recently devoted a special issue (Vol. 50, No. 7, July 2004) to the interface between IT and CS, focusing on problems on sequences and covering topics exemplified by the following list:

- Analysis of algorithms with applications in data compression, prediction, pseudorandom number generation, and classification.
- Computational or descriptive complexity applied to examine the complexity and performance tradeoff in lossless/lossy data compression.
- CS data structures (grammars, trees, and graphs) applied to design codes for sequences.
- Exact and approximate pattern matching techniques applied to context modeling/coding.

- Entropy, mutual information, and information distance concepts, applied to problems of computational biology.
- IT concepts applied to provide bounds in computational learning theory.
- Applications of IT to computational linguistics.
- Applications of IT to information retrieval (especially data mining).
- Analytic information theory techniques for investigating complexity/performance tradeoffs in lossless/lossy data compression.
- Second and higher order analysis of information-theoretic algorithms.
- Information-theoretic bounds for algorithms.
- Applications of information theory to problems of computational molecular biology (e.g., longest common subsequence problem, shortest common superstring problem, genome entropy estimation problem) and security (e.g., intrusion detection).

Because of, among other things, constraints on time to publication, the special issue intentionally left out the interplay of IT and CS in other areas, particularly the interplay in coding, and in networking (including wireless networks, ad hoc networks, and sensor networks). science. Along with interactions between IT and CS in problems on sequences, note that techniques and design paradigms in coding theory and in theoretical computer science have flowed in each direction. Theoretical computer scientists have been applying their knowledge and skill gained in CS in the field of coding theory (particularly in the algorithmic aspects of coding theory) and using error-correcting codes as tools to solve problems in theoretical computer science since the late 1980s. On the other hand, since Shannon's ground breaking work on channel capacity in 1948, information theorists (coding theorists in particular) have been investigating the computational aspects of error-correcting codes, trying to develop either faster decoding algorithms with low computational complexity for existing channel (error-correcting) codes, or new capacity-approaching channel codes with low encoding and decoding complexity. The interplay of IT and CS in the area of coding theory brought about huge successes, as exemplified in the work of Gallager, Berrou, Glavieux, Luby, Mitzenmacher, Shokrollahi, Spielman, Sipser, Richardson, Urbanke, Forney, Sudan, and many others in turbo codes, low density parity check codes (LDPC), graph-based code design, iterative decoding, list decoding and their analysis, some of which have been adopted in the third generation mobile system (3G).

New networking paradigms such as wireless networks, sensor networks, and ad hoc networks do provide opportunities for the interplay of IT and CS in networking. Innovation and outside-the-box thinking in new network paradigms are possible when both the IT and CS/networking communities fully work together.

The past success and huge potential technological and economic benefits in the interplay of IT and CS motivated NSF and us to organize a NSF sponsored workshop—the NSF workshop on Information Theory and Computer Science Interface—to understand challenges lying ahead at the intersection of these two different areas, as well as in the cross-fertilization of these disciplines.

4 The Workshop

The NSF workshop on Information Theory and Computer Science Interface was held in Chicago, Illinois, during October 17 to 18, 2003. This is perhaps the first NSF sponsored workshop addressing fundamental research issues and research training in the interface of information theory (IT) and computer science (CS) even though the interplay between IT and CS dates back to the founding father of information theory, Claude E. Shannon. The objectives of the workshop were to:

- provide a contextual basis for understanding the interplay between information theory and algorithmic aspects of computer science;
- assess the potential impact of computer science and information theory on emerging areas and technology such as bio-informatics and multimedia informatics, and vice versa;
- identify and characterize the major research issues associated with the design, development, and deployment of information systems and algorithms behind them;
- develop a conceptual framework for future cross-cutting research activities;
- examine training of highly qualified personnel in the interplay of IT and CS including graduate student training, undergraduate student training, and their corresponding curriculum developments;

Finally, we were to make recommendations to NSF on how to better utilize NSF's limited resources to promote excellence in research, research training, and research leadership, in the interplay of IT and CS.

Twenty two (22) leading researchers and innovators from both academia and industry participated in the workshop. The organizing committee set the agenda of the workshop. Before the workshop, the participants were asked to submit short summaries describing their recommended priorities for research in the interplay of IT and CS. Summaries received by the organizers were then distributed to every participant for discussion at the workshop. As usual, the workshop adopted the standard format of a single stream of sessions. The first day of the workshop started with the welcome remarks by the organizing committee followed by sessions on the following topics:

- Interplay of IT and CS in networking, especially in wireless networks, sensor networks, and ad hoc networks.
- Interplay of IT and CS in Shannon theory.
- Interplay of IT and CS in communications and coding.
- Interplay of IT and CS in theoretical computer science (TSC).
- Interplay of IT and CS in compression.
- New directions.

Each session was assigned a discussion leader who was responsible for collecting feedbacks and summarizing the discussion outcomes of the corresponding session. The next morning all participants met again to discuss conclusions and recommendations from all sessions and plan the preparation of the report.

The rest of the report reflects the agenda of the workshop, and contains detailed discussion topics, conclusions and recommendations.

4.1 The Interplay of IT and CS in Networking

New networking paradigms such as wireless networks, sensor networks, and ad hoc networks provide excellent opportunities for the interplay of IT and CS to create potentially ground breaking innovation for future communication networks. Traditional layered structures and network architectures for information transport are not justified theoretically. For instance, a recent development in IT—network coding—reveals that simply routing or replicating information at intermediate communication nodes is in general not optimal. Information theory can provide strategic guidance for the design of network architecture for information transport in future communication networks such as wireless networks, sensor networks, and ad hoc networks. On the other hand, practical considerations in network protocol developments and network topology will help information theorists work on the right theoretic frameworks. The interplay between IT and CS greatly resembles the connection of theory with practice. Innovation and outside-the-box thinking in new network paradigms are generally possible only when both the IT and CS/networking communities fully work together.

The workshop recommended the following topics for further investigation:

- Network information theory
  - Asymptotic capacities for large scale or massively dense networks
  - Source coding, including multiple description coding
  - Channel coding
  - Network coding
  - Distributed computing
  - Coding for networking robustness and reliability
- New network architecture and topology
  - Soft edges
  - Dynamic models (mobility, changing node characterization, time varying graphs, etc.)
  - Timing/delay considerations
  - Layered models
  - MIMO
- Resource issues
  - Efficient resource allocation
  - Sharing/cooperation
  - Pricing and incentives
- Network security

4.2 The Interplay of IT and CS in Shannon Theory and TCS

The focus of the discussion in this area was on Shannon entropy versus Kolmogorov complexity. The latter is well surveyed in the book by Li and Vitanyi. The earliest CS applications seem to be due to Freivalds and Paul in the 1970s, and the paper by Paul, Seiferas and Simon played a big role in popularizing the method. Its main effect consists of making certain counting arguments transparent. In its most advanced applications, however, the corresponding counting arguments would be so complicated that they would probably never have been found without the additional insight provided by information theory. The key feature of Kolmogorov complexity is that it is not necessary to have any probabilistic randomness at all.

Usefulness of Shannon and Rényi's entropies in CS became quite clear once the CS community adopted the average case complexity as an equal partner of the worst case complexity. The asymptotic equipartition property (AEP) is often used to prove some combinatorial theorems. For example, several algorithms and data structures on strings (e.g., suffix trees, longest repeated substring) are analyzed by appealing to the AEP.

In Shannon theory, the coding complexity of channel/source coding systems is rarely considered. The performance and coding complexity tradeoff of channel/source coding systems is a more important issue. For lossy source coding theory, partial results have been derived. For more sophisticated coding algorithms, the complexity of the algorithm and the performance need to be investigated.

Another problem that has seen a useful contribution of information theory to computer science is that of random number generation via algorithms that generate equiprobable independent bits from a source of randomness. The problem was initially addressed by von Neumann in 1951, and several advances since then (Hoeffding and Simons, Stout and Warren, Peres, Elias) have shown that the entropy rate is an upper bound for the rate at which it is possible to generate random bits from stationary sources and found an optimal random number generator from stationary finites tate, finite-order Markov sources without making use of the Markov source distribution (other than its order). In later work Vembu, Visweswariah, Kulkarni and Verdu have shown the fundamental limits on the rate at which random bits can be generated from an arbitrary source, and have investigated the properties of universal data compressors as random number generators.

The topics in Shannon theory IT and CS interface recommended for further discussion by the workshop are:

- Network information theory

- Distributed data processing and computation
- Quantum communications, quantum computing, and quantum information theory
- Algorithmic complexity and inference
- Universal estimation of information measures (entropy, divergence, mutual information)
- Information and computation in bio-system
- Dynamic information theory

The topics in theoretical computer science and IT interface recommended by the workshop for further investigation include:

- Algorithms
  - New models of input and output
  - New performance parameters such as entropy of input per source symbol
- Complexity
  - Refined measure of hardness
  - Empirical behavior of algorithms
  - Designing provable cryptosystems
  - New hard function
- Cryptography and security
  - Use information theory to detect intruder attacks
  - Information theory and power attacks

4.3 The Interplay of IT and CS in Communications and Coding

Wireless computing/communication is one area where successful collaboration between CS and IT led to new technology and plenty of interesting open problems. In fact, the database community and the communication community work together to create a new industry.

Here we briefly discuss a different application of communication, namely magnetic recording, on which we did not spend too much during the workshop (it was brought to our attention by Bob Calderbank). Clearly, in this area recording channels are communication channels. This is a big industry, with 250 million hard disc drives sold in 2003, and a 60% density increase per year. Standard tools are coding and signal processing. A small improvement in coding efficiency results in dramatic improvement in manufacturing yields. There are key concepts that require joint research of IT and CS community, including security, distributed storage, multiuser techniques in recording, etc.

There are many other topics recommended by the workshop for further investigation:

- Cognitive radios (spectral allocation)
- Over the horizon sensor communications (localization plus collaborative transmitter)
- Information theory of protocols
- Security for wireless communications
- LDPC
  - Rateless LDPC codes
  - Fountain codes
- New models and applications
- Coding and space-time capacity issues

4.4 The Interplay of IT and CS in Compression

Data compression (known in IT as source coding) is probably the best example of the successful interplay between computer science and information theory. Roughly speaking, Shannon and his followers provided the foundation; the contribution of the CS community is to make data compression algorithms more efficient and to bring algorithmic thinking into this area. The annual Data Compression Conference provides a consistent forum of mutual discussion between these two groups.

We now discuss some more specific problems in the interplay between CS and IT. Since sequential lossless data compression is a particular case of a sequential decision problem, beyond the original focus of source coding on data compression, the information theory (IT) research in this discipline has expanded to areas such as prediction. The computer science (CS) community has also attacked this problem, referred to as "on-line algorithms," in the area of competitive optimality (cf., e.g., the book by Borodin and El-Yaniv). One difference between the approaches is that while the competing (reference) class in the IT literature is chosen such that there exist on-line algorithms for which the redundancy (or regret) vanishes asymptotically, the literature in CS focuses on the "competitive ratio." One CS setting that has been adopted in IT is that of prediction with expert advice. We believe that the importance of on-line problems such as prefetching, caching, web caching, etc., will continue to grow. While the CS community is naturally closer to this type of applications, the IT community has much to offer here.

Specifically, for data compression, a related problem is the investigation of reference classes for analyzing the redundancy of lossless codes for individual sequences. A major contribution of the Ziv-Lempel LZ'78 paper, based on their 1976 definition of LZ complexity, is the reduction of the reference class from the Turing machines (used in Kolmogorov complexity) to FSM's (Finite State Machine), demonstrating a practical algorithm for which the individual-sequence redundancy vanishes for this new comparison class. In Rissanen's stochastic complexity, the comparison classes generally also have a finite number of contexts. Now, can we still have vanishing redundancy for richer comparison classes? The CS community can provide examples of such classes, and the grammar-based codes can be a step in this direction.

An important issue is the trade off between compression and complexity. In general, this issue has received more attention in the lossless compression community, which tends to emphasize also algorithmic aspects, than in the rate-distortion community, which is more Shannon-theoretic in nature.

Finally, another good topic is the better understanding of rate distortion problems, in which the encoder exploits relationships among symbols to achieve better compression without loss of important information (semantic data compression).

Some further topics recommended by the workshop for further investigation include:

- IT analysis of CS structures (suffix trees, grammar, etc) – application to compression architecture (coding, pre fetching)
- On-line prediction
- Redundancy with respect to reference classes

- CS competitive ratio for on line problems
- Interaction with learning theory
- New source models (e.g., for biological data)
- ”Semantic” IT, applications to problems in CS
- Rate distortion theory, backing away from RD optimality
- Joint channel/source coding
- Error resilience.

4.5 New Directions

There are new emerging areas of the computer science and information theory interface that are worth pursuing. Here we discuss in some detail novel applications in molecular biology, information embedding and allocation mechanisms.

The interface between information theory, computer science, and biology is flourishing. The recent assembly of the human genome is generating a lot of attention from both communities. For example, in order to analyze biological sequences we need novel models of sources (e.g., for tandem repeats) and re-shape information theory to include time constraints inherently involved in living organisms.

Let us focus on one facet of biology, namely, the methods by which the human brain handles information. No one really knows an overreaching answer to a question of this breadth and importance. However, Toby Berger had some success at both the single neuron level and the neural network level in analyzing some aspects of their behaviors. For example, a partial differential equation can be derived and solved via Laplace and Fourier transformations, which provides a statistical description of the point process of spiking instants of a cortical neuron whose roughly 10,000 synapses are bombarded by similar point processes that, in general, are mutually statistically dependent. At the neural net level one can use a finite-state channel with feedback of its vector output to its vector input to mimic several of the principal characteristics observed within certain real neural nets. At the synapse level it is believed that major roles are played by various ion channels. Finally, the ends to which the brain puts information are at best poorly understood at present and perhaps totally misunderstood. It is our belief that the theories of communications, control, estimation, decision-making, pattern recognition, learning, adaptive systems, and so on, can play a meaningful part in improving our understanding about the short-term and long-term ends that the brain pursues when processing information.

The development and testing of meaningful theories of brain function requires an understanding of how millions, or even tens of millions, of neurons work together. We need to know their electrical and chemical states at a time resolution on the order of a millisecond down to perhaps even a microsecond in compartments, some of which have nanometer dimensions. Since experiments cannot presently inform us at this level of detail, we try to turn to simulation. But most simulations turn unstable once tens of thousands of neurons are included, and hence fall at least two and often three orders of magnitude short of being definitively informative. Computer scientists, information theorists, control theorists, and neuroscientists need to work together to overcome these shortcomings.

The above theme is common to other endeavors, too. Wireless and sensor networks will also soon have hundreds of thousands of nodes, maybe millions, in a few years. Semi-conductor chips already have millions of elements deposited upon with complex connectivity graphs. We need, as in the RNN application, to understand how information flows through structures, how this flow may be limited by the rate of energy dissipation, and what "coding" might be able to improve network performance. Error bounds that give "tight" approximations to (optimum) network behavior for large but still finite sizes are needed; order of decay is key, of course, but not enough in itself; we also need to bound the multiplicative constants inherent in $O(\cdot)$ results.

Another area of emerging interest to CS and IT is information embedding. While the information-theoretic analysis of information embedding has developed mostly as an application of the ideas of coding with side information in Shannon theory, one can think of a parallel with the development of the theory of error-correcting codes 40 years ago, and the influence of the more algorithmic-oriented thinking of computer science. At the time, we saw on one hand a Shannon-theoretic trend, centered around random coding, capacity and error-exponent computation, etc., but also, on the other hand, a "parallel track," developed by information theorists, mathematicians, and computer scientists, in which the problem was detached from its roots and studied in a combinatorial framework, where the target is the minimum-distance vs. rate tradeoff, rather than aiming at channel capacity. It was mostly this combinatorial track that attracted the ideas of algorithmic computer scientists, and so the development of practical error-correcting codes, even codes that are sub-optimal from a Shannon-theoretic perspective, was done in a well-defined theoretical framework, with bounds to be achieved.

Finally, algorithms, complexity and information in the design of allocation mechanisms and games were brought up. Mechanism design deals with allocation of goods or resources among multiple agents. Often the value of a resource to a particular agent is information privately held by that agent. An example is an auction, in which the seller does not know how much the object to be sold is worth to each of the buyers. Another example might be a communication network in which a buyer would like to purchase the use of links to form a path from a source to a destination, but the buyer does not know the cost to the provider of each link. Mechanisms may have certain desirable properties, such as (i) agents to reveal their true valuations, (ii) making socially efficient allocations, (iii) providing agents to participate, or (iv) being budget balanced. Mechanisms can be designed under the assumption that agents will play dominant strategies, or possibly Bayesian Nash equilibrium strategies.

Algorithmic and complexity considerations are prominent in the implementation of mechanisms, because the mechanisms with desirable properties can in some cases be quite complex. A search for low complexity, in some cases distributed, mechanisms has recently been launched within the theory of computing community.

Information considerations are also prominent in the implementation of mechanisms. For one, the assumption of private information on the part of participants is a major feature of the theory of mechanism design. For another, there is often a tradeoff between how well a distributed allocation mechanism works and how much communication it requires, or how much information is available to the agents. Perhaps in the context of mechanism designs or games, a value can be placed on information.

Further topics recommended by the workshop for future investigation include:

- Coding for nano-structures
- Massively dense systems
- IT/CS in biology
  - Databases (genomics, medical information, protienomics)
  - Organism (inter and intra)
  - Computing with biological materials

5 Recommendations to NSF

We recommend a new initiative by NSF on basic research that exploits the interface between Computer Science and Information Theory. The Information Technology revolution that has affected Society and the world so fundamentally over the last few decades is squarely based on computation and communication, the roots of which are respectively Computer Science and Information Theory. These two components of Information Technology have developed in a divergent way that has created almost entirely distinct communities with their own language and culture; these have often been at odds with each other and have in fact contributed to the somewhat chaotic character of the Information Infrastructure. It is high time, as the participants in this workshop unanimously agree, to revisit and exploit the deep connections between these two disciplines. The current organization of the National Science Foundation makes it difficult to identify the programs that could harbor and support research efforts on the foundations and interfaces of these two core disciplines. So, a special initiative by the Foundation would rekindle research on the many maturing topics outlined in this report that point to the strong interplay between Computer Science and Information Theory.