

MOSIS Semiconductor Implementation Service  
Digital Phone Communications Sports Data Analytics

**ISI FOUNDED BY KEITH UNCAPHER**

ICANN SPINOFF PACKET VIDEO ISI Arlington ONLINE PAYMENTS

FASTXchange Spinoff Packet Radio Terminal Concept

USC-LOCKHEED MARTIN QUANTUM COMPUTING CENTER

MOSIS Alliance with Intel Corporation

Human Trafficking Search Tool ARPANET INTERFACE

RHETORICAL STRUCTURE THEORY



CELEBRATING 50 YEARS OF INNOVATION

2022 ANNUAL REPORT

Internet Society  
Social Media Bot Detection

DETERLAB CYBERSECURITY TESTBED

Practical Rhetorical Theory Internet2 Collaboration

ISI Boston ISI 50th Anniversary

USC Institute for Creative Technologies Spinoff

Internet Domain Name System

SoCal Seismic Hazards Map

USC Viterbi School of Engineering

ARPANET Voice over IP

ARPANET Addressing and Request for Comments

Nano-Satellite  
Reconfigurable Robots

# *Information Sciences Institute*

is a world leader in research and development of advanced information processing, computing, and communications technologies.



## **CALIFORNIA**

4676 Admiralty Way #1001  
Marina del Rey, CA 90292  
310.822.1511



## **MASSACHUSETTS**

890 Winter Street #115  
Waltham, MA 02451  
781.622.9790




## **VIRGINIA**

3811 Fairfax Drive #200  
Arlington, VA 22203  
703.812.3700



## **REMOTE RESEARCHERS**



<b>2</b>	Welcome to ISI
<b>4</b>	ISI 50 <sup>th</sup> Anniversary
<b>8</b>	ISI History
<b>18</b>	Symposium on the Future of Computing Research
<b>19</b>	Leadership
<b>20</b>	ISI in Numbers
<b>21</b>	Scientific and Technical Portfolio
<b>22</b>	Diversity, Equity, and Inclusion
<b>23</b>	In Memoriam
<b>24</b>	Divisions Artificial Intelligence Division Computational Systems and Technology Division Informatics Systems Research Division Networking and Cybersecurity Division The MOSIS Service
<b>29</b>	Centers and Labs
<b>34</b>	Keston and ISI Exploratory Research Awards
<b>36</b>	Research Highlights
<b>53</b>	Viterbi Data Science Program
<b>54</b>	New Sponsored Research Awards
<b>58</b>	Honors and Awards
<b>62</b>	New Research Directors, Team Leaders, and Leads
<b>63</b>	New Faculty Appointments and Promotions
<b>64</b>	2022 PhD Graduates
<b>65</b>	Visitors and Postdoctoral Scholars
<b>66</b>	PhD Students
<b>69</b>	Student Research Programs
<b>70</b>	Publications



This year, we remembered the past and looked to the future. September’s celebration of ISI’s 50 years of information science research was terrific. Not only did we review our research accomplishments, we also celebrated the ISI community. Immediately after that celebration, we looked forward—sponsoring a US-wide seminar on the future of computing research.

We worked hard this year to maintain and support our ISI community as COVID morphed from pandemic to endemic. ISI now offers extreme flexibility to staff, to work at home or in the office, or some of both, according to individual and team needs and preferences.

We are also working hard to promote diversity, equity, and inclusion, not only within ISI itself, but also with our research topics and outreach programs.


ISI researchers are driving novel advances in computer vision, language understanding, cybersecurity, quantum computing, electronics, networking, and other fields. Nineteen ISI graduate students received their doctorates this year and accepted positions at leading technology companies, including Apple, Google, Facebook, and Amazon Alexa AI.

This report presents Information Sciences Institute in a package. You’ll see our numbers, our people, and a sampling of our wide-ranging research, from AI that studies gender biases in science to demonstrating a quantum technology that is an “optical correlator” to consolidate data transmission and data processing. You’ll see the list of our many publications. But as you go through the report, I hope you will also get a feel for our organization and our energy.

We thank our sponsors and our collaborators for their trust in our work. And we look forward to more scientific and engineering progress, more teamwork, and more contributions to the computing universe.

A handwritten signature in black ink, appearing to read 'Craig A. Knoblock'.

Craig A. Knoblock  
Keston Executive Director  
USC Information Sciences Institute





## Marina del Rey, California

ISI occupies nine floors of the Marina Towers South building, west of Los Angeles. About 325 researchers, staff, graduate research assistants, faculty, and students are based here, working in all areas of computer science and engineering. The institute’s administrative departments are also located here. The Space Engineering Research Center and the USC-Lockheed Martin Quantum Computing Center are both next door.

Free shuttle service runs between ISI, the USC Institute for Creative Technologies, and the USC University Park campus. ISI is a 20-minute ride from Los Angeles International Airport; it is accessible by public transit and the 22-mile coastal bike trail.



## Arlington, Virginia

ISI Arlington is located in Northern Virginia, near Washington, D.C. Three Arlington researchers hold appointments as research professors in the USC Ming Hsieh Department of Electrical and Computing Engineering. ISI’s Secure and Robust Electronics Center and Application Specific Intelligent Computing Lab are both located here. In addition to collaborating with researchers in the other ISI locations, researchers at ISI Arlington collaborate with USC academic departments and centers, other universities, commercial companies, and defense contractors. Student interns are recruited from universities around the world to work at ISI Arlington, and from local and regional universities.



## Boston, Massachusetts

Located in the Greater Boston area, the Waltham location celebrated its fifth anniversary this year. Major research areas at ISI Boston include artificial intelligence, deep learning, multimedia and natural language understanding, and quantum information science. Researchers lead many projects in collaboration with universities and industry. ISI Boston also partners with local universities and universities around the country to host interns and co-op students. ISI’s Laboratory for Quantum-Limited Information, which is dedicated to understanding and demonstrating the fundamental physical limits for extracting information from physical signals, is located here.

## 50TH ANNIVERSARY GALA

ISI staff and alums celebrated ISI's 50th anniversary in September with dinner, talks, live music, and the documentary *Cloudwalkers: ISI and the Inventors of the Future*.



Keston Executive Director Craig Knoblock, USC President Carol Folt, and Engineering Dean Yannis Yortsos highlighted ISI's contributions to network and computing technology and to USC. They were introduced by ISI Senior Director Yigal Arens. (left to right)



Keston Executive Director Craig Knoblock

ISI's Keston Executive Director Craig Knoblock gave a brief history of ISI and presented his ambitions for ISI. "For the next 50 years, I want imagination to continue to reign at ISI. Our research must expand further and become even more collaborative with other disciplines. Think about how we could contribute to the medical field, to environmental research, and more!"





USC President Carol Folt

USC President Carol Folt’s talk cited ISI as holding a “central place at Viterbi.” She noted that “about 50 percent of Viterbi’s research funding comes through ISI, and for 50 years, you’ve been mentoring computer scientists and innovators.”



Yannis Yortsos, Dean of the USC Viterbi School of Engineering

Yannis Yortsos, Dean of the USC Viterbi School of Engineering, said, “ISI’s birth almost immediately transformed USC engineering from good to great in the realm of computing. For five decades, ISI has revolutionized computing worldwide. I believe the Institute’s best days lie ahead.”

“Clearly, the relationship between ISI and USC Viterbi has contributed greatly to our goal of leveraging technology for the benefit of humanity.” He said, “I cannot wait to see all the transformative discoveries borne from this singular relationship over the next 50 years.”



Amanda Keston and Daniel Keston, son of Michael Keston. Michael Keston created the Michael and Linda Executive Directorship Endowment, ISI’s first philanthropic gift, in 2015.



## 50TH ANNIVERSARY GALA (CONTINUED)



USC President Carol Folt chatted with Bill Swartout, who joined ISI in the early 1980s. Bill is now the Chief Technology Officer at the USC Institute for Creative Technologies, an ISI spin-off.



Engineering Dean Yannis Yortsos (left) and Keston Executive Director Craig Knoblock



Bob Balzer (left) was one of the original researchers in 1972 when Keith Uncapher started ISI. Next to Bob is ISI Senior Director Yigal Arens.



Steve Crocker (right) joined ISI in 1974 as a senior researcher. He later served as chair of ICANN, the Internet Corporation for Assigned Names and Numbers. Steve won the IEEE Internet Award in 2002. On the left is Carl Kesselman, Director of ISI's Informatics Systems Research Division, who has received the 2023 IEEE Internet Award.



Both Eve Schooler (left) and Bob Felderman are ISI alums. Eve was awarded the IEEE Internet Award, along with ISI alum Steve Casner, in 2020.



Rodney Van Meter, an ISI alum, is Vice Center Chair of the Keio Quantum Computer Center. He's catching up with Dave Bilkis, another ISI alum. (left to right)





Paul Mockapetris, who developed the Internet Domain Name System while at ISI, contributed to the making of the ISI documentary *Cloudwalkers: ISI and the Inventors of the Future*. Paul received the IEEE Internet Award in 2003. The documentary was written and directed by Daniel Druhora. (left to right)



The documentary *Cloudwalkers: ISI and the Inventors of the Future* begins with the founding of ISI and describes its role in designing, developing, and running the Internet. It delves into the advancements in cybersecurity and cyber defense made by researchers at ISI and others. It offers glimpses of some of the surprising work being done at ISI today. Watch *Cloudwalkers* trailer.



### Founding ISI

In 1972, **Keith Uncapher** headed RAND Corporation's computer sciences department but was facing budget issues. He approached UCLA about creating a computer science institute and then reached out to George Bekey, Chair of the USC Electrical Engineering Department. Dr. Bekey quickly connected him with Zohrab Kaprielian, Dean of Engineering and USC Executive Vice President. Dean Kaprielian moved fast, and the Information Sciences Institute was launched five days later. DARPA funded the institute, then just Keith Uncapher and a handful of researchers. By the end of the first year, though, ISI had 61 people, not only researchers (many from RAND) but also consultants, support staff, research assistants, and student aides.



### Why *information sciences*, not *computer sciences*?

In a 1989 interview, Keith Uncapher said:

*Here, in the beginning of 1972, the major computer science departments were still largely involved in traditional forms of computer science ... no intent to be application-oriented, or try to discover science required for major applications, ... exploratory systems or proof-of-concept kind of systems. So the combination of leaving RAND and having the high probability of some very good people available, I went to DARPA and said, "I would like to start a new venture. It makes sense to me. It would be off-campus, but would be totally a part of the university ... I proposed 40% basic research and 60% applications focused on problems and opportunities important to the Department of Defense."*

### ARPANET Growth

One goal of the new institute was to establish ISI as a major host system on the ARPANET, the forebearer of the Internet. ISI designed and produced its host interface for ARPANET. By the end of the first year, ISI operated the busiest node on the ARPANET, with some 400 external users. Over the following decades, ISI played a pivotal role in creating and managing the ARPANET, then the Internet, including core concepts, technical standards, and ongoing functionality.



### ARPANET Voice over IP (VoIP)




Danny Cohen implemented online voice communications, the first use of network packet switching for real-time voice applications.




### First Portable Remote-Access Terminal

ISI developed the first portable remote-access terminal. A prototype terminal was delivered to DARPA (then ARPA) in July 1973. The ISI technical report noted that it weighed 20 pounds and could fit under an airplane seat. The director of DARPA was an early adopter.



## ARPANET Addressing and Request for Comments

Jon Postel moved from UCLA to ISI, where he continued managing global address allocation and maintaining the Request for Comments (RFC) series, the cornerstone of the Internet standards process. Jon continued in this role until his sudden death in 1998.



## Packet Radio Terminal Concept

The notion of a packet radio terminal was conceived at ISI. A forerunner of the modern smartphone, the system proposed communicating by radio with a keyboard and display screen.

## Video Display

ISI modified a Hewlett-Packard terminal to create the first full-page, character-based display—a prototype for early PC workstation terminals until the arrival of LCD in the mid-90s.

## Scholarship

The 1974 ISI Annual Technical Report noted 7 PhD theses in progress.

## MOSIS: Metal Oxide Semiconductor Implementation Service

Led by Danny Cohen, the MOSIS service began operation in 1981. This was “computerized commerce,” in Keith Uncapher’s words, with transactions conducted via network. MOSIS slashed chip production costs by consolidating multiple customers’ ICs onto shared wafers.



## FAX via Network

ISI’s hardware lab designed a FAX interface to the ARPANET. Instead of sending the data over a phone line, the data was digitized and incorporated into a file that could be stored, printed, or sent over the ARPANET on demand. Packet-switched network communications research and development was a major thrust of the new institute. This included voice, data, and image use in remote conferencing applications.



*Craig Rogers, who started programming at age 12, came to work at ISI in 1980. Craig was working at UCLA and had collaborated with some ISI people there. 42 years later, when asked why he left UCLA for ISI, Craig said, “Because ISI had people like me and the newest toys.”*

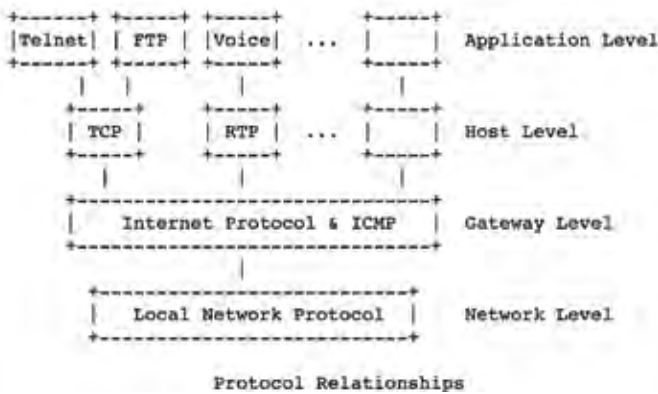


ISI continued to advance network technology, with breakthroughs in artificial intelligence research and additional communications modes through packet technology. In 1988, **Herb Schorr**, IBM's

vice president for research, became executive director of ISI. He broadened ISI's core strengths—networking and software engineering—to diverse technologies, recruited new researchers, and encouraged cross-disciplinary and cross-university research. Funding sources grew to over 20 federal agencies in science, medicine, space, and defense, plus corporate sources.

### TCP/IP Network Protocol

Jon Postel and ISI colleagues made major contributions to the Transmission Control Protocol/Internet Protocol (TCP/IP) network protocol in RFC 793, laying the foundation for the modern Internet.



*FAX, phone, whatever . . . As Bob Parker, who headed ISI's Hardware Development Lab, says, "The power of packet switching was that everything was simply data: voice, video, FAX streams, it was all the same."*

### Digital Phone Communication

ISI engineers created an interface enabling phone calls to travel over ARPANET for the first time. The interface linked the traditional phone system and ARPANET by converting phone signals from analog to digital and back. Digital transmission became the communications industry standard.



### Soar AI Architecture

Soar was a groundbreaking cognitive architecture—a hypothesis concerning the fixed structures and processes that yield a mind. Soar has been applied across a vast range of tasks to model human cognition and support the development of complex artificial intelligent systems. It has been used to simulate virtual humans supporting face-to-face dialogues and collaboration within a virtual world. Paul Rosenbloom co-led the effort from 1983 to 1998, first at CMU and then at ISI. Soar is still supported and extended at the University of Michigan and the Center for Integrated Cognition. At ISI, one application of Soar in the '90s was the simulation of helicopter pilots and commanders for large-scale distributed simulation. In a retrospective, an ISI colleague wrote about "...using it to fly helicopters in teams that a commander found to be indistinguishable from human pilots. I thought that was remarkable."



## Internet Domain Name System (DNS)

ISI researcher Paul Mockapetris invented the Internet Domain Name System (.edu, .net, .com, ...). DNS automatically translates text addresses to numerical addresses that computers can understand. In an interview on NPR 26 years later, he was asked about the first domain name that he made. He replied, “Well, I was doing the work at ISI, so I’m sure it was isi.edu.”




## Packet Video

ISI created the specification for packet video in 1981. Just a few years later, ISI debuted packet video and transmitted via a satellite on the building’s roof. (This image is the actual satellite dish.) Video is now a fixture of the contemporary Internet.



## Network Growth


Bob Braden moved from UCLA to ISI. He contributed to fundamental Internet communications protocols, operating standards, and architecture and co-edited the Request for Comments (RFC) document series.



## Rhetorical Structure Theory

AI researchers invented Rhetorical Structure Theory (RST), which explains precisely what makes written texts coherent: the relations between parts of text. RST became the basis for a body of discourse-related natural language research. Natural language research became one of ISI’s major areas.

### Generic RST Schema




## Internet Society

The Internet Society was founded to promote the Internet’s evolution—including its standards, administration, and technical infrastructure. Jon Postel raced a colleague to become its first member. The Internet Society is now a global advocacy nonprofit with 128 chapters.



## Multicast Online Transmission

Multicast transmission over the Internet, or Mbone, was co-created by ISI researcher Stephen Casner. Two years later, Mbone transmitted a Rolling Stones concert worldwide via the Internet for the first time.





By ISI's third decade, the Internet had gone mainstream. Personal computers offered online bulletin board systems (BBSs), dial-up Internet service, and computer games to regular people. ISI went bicoastal, opening an office in Arlington, Virginia, near Washington, D.C. Some of ISI's contributions include moving money by network, coordinating big data and big computers, the first step toward cloud computing, and advances in Internet management. In AI, researchers advanced the field of information integration, the foundation of much of today's computing. And what about Y2K? Would every computer and network come crashing down the minute 1999 ended? ISI's networking division director and some colleagues went to ISI to wait and watch the clock. It was a quiet evening.

Later that year, Director Herb Schorr echoed Keith Uncapher in an interview with the Los Angeles Business Journal.

**Question:** *For those who've never heard of it, describe ISI.*

**Answer:** *We build systems, in contrast to most universities. We're sort of between a corporate laboratory and a university in the sense that we do good research by building something new.*

### Online Payments

ISI researchers devised NetCash, followed by NetCheque clearing and conversion one year later.

NetCash and NetCheque were among the first Internet-based electronic payment systems.



### Internet Commerce

In a 2011 interview, Danny Cohen said: "... then one day I said that the mechanisms that we have to run MOSIS could be used to run any general electronic commerce." He and his team developed the FASTXchange online brokering system for aerospace, defense, and electronics parts. This pioneering software solution was eventually commercialized in 1997.



### More Packet Communication

ISI's Eve Schooler and Steve Casner developed standards for Internet multimedia, making formative contributions to the design of Internet multimedia protocols.

### Grid Computing



ISI researchers, working with Argonne National Laboratory collaborators, developed grid computing to connect computing resources and provide massive amounts of computing power and data for scientific research.

Grid computing architectures were utilized to perform the computing for three Nobel Prizes. Grid computing laid the groundwork for cloud computing.

## Robotic Soccer ... and more



ISI's robot "Dream Team" became the world champion in mid-sized robot soccer, tying with a Japanese team for overall honors and outperforming 40 teams.

Work in robotics continued,

and in 2008, an ISI robotics team deployed self-configuring polymorphic robots to win the International Conference on Robotics and Automation contingency competition.

## ICANN

As Internet use expanded, ISI's Jon Postel and collaborators worked with the White House and Department of Commerce to create the Internet Corporation for Assigned Names and Numbers (ICANN).

ICANN was based in the same Marina del Rey building as ISI.




## Virtual Reality

Some of ISI's virtual reality researchers moved to the newly created USC Institute for Creative Technologies (ICT), an ISI spin-off. ICT continues to pioneer virtual reality and other immersive technologies for military and civilian uses.



## Knowledge Representation: Loom




The Loom project developed advanced tools for knowledge representation and reasoning in AI. Loom, and now PowerLoom,

is a language and environment for constructing intelligent applications. Loom has been distributed to over 80 universities and corporations and is used in numerous DARPA-sponsored projects in planning, software engineering, and intelligent information integration.

## Adaptive Computing ICs

ISI accelerated high-performance computing by utilizing FPGAs, integrated circuits that can be configured after manufacturing. The chips were deployed in a novel architecture for embedded computing and, for the first time, in cluster computing.



## Internet Speed

ISI researchers collaborated on the development of Internet2, which set a land speed record.



## Practical Rhetorical Structure Theory



ISI research demonstrated how Rhetorical Structure Theory (RST) could be used in practical discourse parsing and text summarization applications. Daniel Marcu implemented machine learning models that understand the mechanics and rhetoric of well-written text. He worked with the Educational Testing Service to deploy RST-based discourse parsing for scoring millions of essays automatically each year. Language Weaver, the first phrase-based statistical machine translation software company, was later spun out of ISI (years ahead of Google and Microsoft).

## Electric Elves

AI researchers built a system of software agents that helped ISI staff reschedule meetings, select presenters, track locations, organize lunch meetings, and other tasks. The system was tied to workstations, fax machines, and mobile devices. Each proxy was called Friday (after Robinson Crusoe's manservant) and acted on behalf of its user in the agent team. The system ran 24/7 for months.



During this decade, Herb Schorr's aim of broadening ISI's strengths to diverse technologies took hold. ISI became home to the USC Space Engineering Research Center and won NASA awards for developing space technologies. Grid computing expanded to new fields—and went commercial. ISI's artificial intelligence group grew, and ISI established a major cybersecurity testbed. The Department of Energy became a major sponsor. Then a quantum computer system arrived at ISI, in the USC-Lockheed Martin Quantum Computing Center.



### NASA Space Technology Research Projects

In 2004, ISI won three NASA grants totaling over \$58 for developing new space technologies. ISI received roughly half the funding, with subcontractors including other universities and aerospace companies. The three projects advanced the development of mechanisms that can work with minimal human control in space:

- *Modular, Multifunctional Reconfigurable* robotic system for construction and maintenance
- *Coordinated Multisource Maintenance-on-Demand* space logistics system
- *Fault-aware, Modular, Reconfigurable Space Processor* control system

The projects capitalized on ISI's strengths in AI, automated construction, sensor and machine networks, and distributed collaborated systems.

### Robotic Building Construction (“House Printing”)

Robotic construction methods were pioneered at ISI in a technique known as Contour Crafting, supported by the National Science Foundation. Contour Crafting, essentially large-scale 3D printing, builds up shapes in layers by controlling the flow of liquid building materials using two movable, programmable trowel-like tools deployed around a nozzle.



### DETER Cybersecurity Testbed



ISI partnered with UC Berkeley on Cyber Defense Technology Experimental Research (DETER), a major cybersecurity testbed. The DETERlab provides a “mini-Internet,” a safe place for thousands of researchers to test

cyberattacks and defenses. DETER later added “research on research” capability and now serves 600+ cybersecurity experimenters in 16 countries. The 805 square foot lab is located on the 11th floor of ISI's Marina del Rey location. It is the world's largest public testbed for cybersecurity.

### Virtual Learning Games



The Alelo company originated from an ISI research project, *Tactical Language Trainer*, which developed an interactive language training system employing an intelligent agent coach and gaming techniques.



## Space Engineering Research Center at ISI



USC astronautical engineering professors started the Space Engineering Research Center (SERC) in 2006. Herb Schorr, ISI's director, provided lab space at ISI. Two of the ISI researchers who joined the center eventually became associate directors of SERC. Graduate and undergrad

students get hands-on experience with spacecraft technology at SERC and are directly involved in the research and development of space hardware and software.

## Department of Energy Research

In 2006, ISI played a significant role in supercomputing research sponsored by the Department of Energy. The Scientific Discovery Through Advanced Computing program, SciDAC, funded three ISI projects:

- *Performance Engineering Research* enhancing the performance of SciDAC applications on petascale systems
- *Getting the Science out of the Data* scientific data management to help scientists spend more time studying their results and less time managing their data
- *Sharing a World of Data* scaling up the Earth System Grid to manage petascale data to enable faster, easier sharing of climate change research

## Recognition for ISI's RFC Stewardship

A coalition of Internet organizations, including the Internet Society, the Internet Engineering Task Force, the Internet Architecture Board, and the Internet Research Task Force, sponsored an event in 2008 to recognize ISI's years as editor of the Request for Comments (RFC) series. RFCs are the documents that describe and shape the workings of the Internet. Over 50 computer pioneers and leaders key to the Internet's development attended, including Steve Crocker, the author of RFC 1.



## Language Processing: "What Is Artificial Intelligence?"

Three senior ISI researchers were among the faculty members from eight universities who collaborated with IBM to design learning systems for the Watson computer. Watson famously beat humans on the Jeopardy! TV quiz show.



## How Big is the Internet?

ISI researchers mapped the entire Internet address space, comprising 4.3 billion Internet addresses.

## De-encrypting a Manuscript

An ISI natural language processing researcher, with two Swedish collaborators, decoded the famed Copiale Cipher, a hand-lettered, 105-page manuscript that had resisted cracking since the 1730s. The manuscript was discovered in an academic archive in the former East Germany and contained Greek letters, Roman letters, and abstract symbols. The manuscript revealed the rituals and political thoughts of an 18th century secret society in Germany.



## Quantum Computer System

ISI became the home of the 10 million dollar D-Wave One Adiabatic Computer, the first in operation outside D-Wave facilities. The USC-Lockheed Martin Quantum Computing Center (QCC) is a joint effort of Lockheed Martin and USC, where faculty, students, and researchers conduct basic and applied research. In May 2022, the system was upgraded to the fifth-generation Advantage Quantum, one of the first in the US.




## ISI HISTORY: THE FIFTH DECADE

ISI's fifth decade began with a new director, **Prem Natarajan**. He was an executive at Raytheon BBN and an expert in optical character recognition, speech processing, and multimedia analysis. Under his tenure, broader, larger, more collaborative research projects were launched. The new VIMAL computer vision lab started up and received substantial funding. In 2015, ISI received a generous



gift from Michael and Linda Keston that endowed ISI's directorship and continues to fund exploratory research. In 2018, **Craig Knoblock**, a longtime ISI artificial intelligence researcher, was named ISI's Keston Executive

Director. ISI projects range from societal issues related to technology to quantum computing to small satellites.



### Quantum Confirmation

Quantum researcher Federico Spedalieri of ISI witnessed the "entanglement effect," demonstrating that the D-Wave system is actually performing quantum processing.



### Sports Analytics

Two ISI researchers applied data analytics to basketball and successfully launched a startup that commercializes their approach. Second Spectrum is now the dominant sports analytic company for the NBA and soccer's English Premier League.



### Fighting Modern Slavery



ISI researchers developed an Internet search tool to expose and combat human trafficking. Their cloud-based analytics

tool, Domain-specific Insight Graphs, uses open-source software to identify and extract web information, producing meaningful, searchable information for law enforcement.

### Detecting Gravitational Waves

Pegasus software, which automates scientific computational work as workflows, was credited with facilitating MIT and Caltech's Nobel Prize-winning detection of gravitational waves. The discovery, predicted by Einstein's Theory of Relativity and made by Laser Interferometer Gravitational-Wave Observatory (LIGO) physicists, is among the scientific projects managed with Pegasus.



### Online Social Behavior

ISI's research contributed to online social behavior research. Among them: social network illusions that cause people to believe rare events are common, why more content doesn't go viral, and how Twitter may track disasters more quickly than the US Federal Emergency Management Agency.

## Bots . . . and Elections

An ISI team created the nation’s fastest, most effective method to identify fake Twitter accounts, or bots. In later work, Emilio Ferrara’s work about the possible consequences of fake Twitter bots on the 2016 election received much attention. By leveraging state-of-the-art bot detection algorithms, Ferrara and his team discovered that pro-Donald Trump and pro-Hillary Clinton bots created a surprisingly high percentage of the political discussion on Twitter to distort the online discussion regarding the elections. This work generated more than 40 news stories in international media, including the cover of *MIT Technology Review*.




## ISI Boston

ISI established a third research location in the Greater Boston metropolitan area, in Waltham, Massachusetts.



## A Computer’s Guide to Humans



Fifteen years in the making, *A Formal Theory of Commonsense Psychology: How People Think People Think*, dubbed a “Computer’s Guide to Humans,” was published by Cambridge University Press. The ISI and ICT authors wrote this book to explore human commonsense reasoning and serve as a foundation for developing human-like artificial intelligence.

## USC’s Third Satellite

Dodona was one of the 80+ satellites launched on SpaceX’s Transporter 3 “rideshare launch” in January 2022. The Dodona nanosatellite is a project of the Space Engineering Research Center (SERC). Dodona was deployed to fly new payloads and demonstrate various control and communications techniques.



Screenshot from January 13th SpaceX launch. Transporter 3 is carrying USC satellite Dodona.

The Dodona project demonstrated: a) a new B-Dot control algorithm for optimal stability control with low power using torque rods; b) a rotating multi-packetized beacon transmission source to prove backup to the primary communications system; and c) data store-and-forward software for non-consistent ground contacts. It uses the flower petal configuration for primary solar array power.



Where is Dodona?

*“We are now beginning to witness that ISI and USC are becoming an attractive force to the U.S. computer science community.” Keith Uncapher wrote in 1973. “I believe that we will be very successful in attracting very, very good computer people.”*

# SYMPOSIUM ON THE FUTURE OF COMPUTING RESEARCH

SEPTEMBER 12-13, 2022 IN-PERSON AND ONLINE

HOME ABOUT PROGRAM WHY ATTEND? SESSION VIDEOS

SPONSORED BY  
**USC Viterbi**  
School of Engineering



On September 12 and 13, ISI hosted the Symposium on the Future of Computing Research. Recognizing the growing power of computing research as computers and computing become embedded in society, ISI brought together thought leaders and visionaries to consider the new goals, constraints, and challenges in the field of computing research.

The symposium featured interactive discussion sessions, talks, a roundtable with early career researchers, and informal, self-organizing discussions. It was held simultaneously online and across multiple physical locations in the US.

**Dr. Sethuraman Panchanathan**, Director of the US National Science Foundation, came to ISI headquarters to give the first day’s keynote presentation, *Innovation Anywhere, Opportunity Everywhere: Accelerating the Frontiers of Computing Research*.

Panels and discussions led by professors from leading universities featured *Identifying and Evangelizing Emerging Technical Directions*, *Recruiting Talent and Fostering Research Careers*, and a *Student & Early Career Researcher* roundtable.

The second day began with a “fireside chat”—an interview with **Eric Horvitz**, Chief Scientific Officer at Microsoft, by ISI’s Terry Benzel and John Wroclawski. Dr. Horvitz responded to questions from remote symposium participants as they came in.

David Tennenhouse of the Carnegie Endowment moderated the *New Models for Research and Research Support* discussion.



View the symposium talks.



## Executive Leadership

**Craig Knoblock, PhD**  
*Keston Executive Director*

**Terry Benzel, MA, MBA**  
*Associate Director*

**Stephen Crago, PhD**  
*Associate Director*

**James Whalen, MPhil**  
*Associate Director*  
*Chief Financial Officer*

**Eileen Lu, BS**  
*Chief Information Officer*  
*Director, Computing and Information Services*

**Yigal Arens, PhD**  
*Senior Director for Administrative Affairs*

**John Wroclawski, MSEE**  
*Senior Director for Strategic Initiatives*

**Lifu Chang, PhD**  
*Director, The MOSIS Service*

**Elizabeth Boschee, BS**  
*Director, ISI Boston*

**Yolanda Gil, PhD**  
*Viterbi Director for Major Strategic AI  
and Data Science Initiatives*

**Derek Mikuriya, MBA**  
*Director, Human Resources*

---

## Research Leadership

**Yigal Arens, PhD**  
*Acting Director, Artificial Intelligence Division*  
*Director, Emerging Activities Group*


**Elizabeth Boschee, BS**  
*Associate Director, Artificial Intelligence Division*

**Stephen Crago, PhD**  
*Director, Computational Systems and Technology Division*

**Carl Kesselman, PhD**  
*Director, Informatics Systems Research Division*

**Terry Benzel, MA, MBA**  
*Director, Networking and Cybersecurity Division*

**David Balenson, MS**  
*Associate Director, Networking and Cybersecurity Division*



STAFF, FACULTY, AND STUDENTS

440+

FACULTY

35

- Astronautical Engineering
- Biomedical Sciences
- Civil and Environmental Engineering
- Communication
- Computer Science
- Electrical and Computer Engineering
- Industrial and Systems Engineering
- Physics and Astronomy
- Population and Public Health Sciences
- Preventive Medicine
- Spatial Sciences

QUANTUM COMPUTERS

1

2022 SEMINARS

64

POSTDOCS

7

STAFF AND FACULTY

7

MBA degrees

53

master's degrees

65

bachelor's degrees

93

PhD degrees

PHD STUDENTS

110

2022 NEW RESEARCH GRANTS

53

2022 RESEARCH EXPENDITURES

\$72 million

SUMMER INTERNS

23

TOTAL OFFICE SPACE

160,196 square feet

ISI LOCATIONS



Marina del Rey, CA



Arlington, VA



Waltham, MA

**Novel Electronics Technologies and Architectures**

- Integrated circuit fabrication brokerage
- Multi-project integrated circuit wafers
- State-of-the-art semiconductor process technologies
- Integrated circuit design automation

**Space Engineering Research**

**Informatics Systems Research**

- Sociotechnical platforms for data-centric discovery
- Informatics cyberinfrastructure
- Collaboration infrastructure for biomedical research

**High Performance, Edge, and Quantum Computing**


- Secure, robust, and trusted electronics
- Heterogeneous and cloud computing
- High performance computing architectures and software
- Quantum information science; quantum computing
- Science automation technologies

**Computer Networks and Cybersecurity**

- Network and security measurement, analysis, and defense
- Network infrastructure supporting science and applications
- Research, methods, and infrastructure for cyber experimentation
- Social engineering attacks
- Binary program analysis, vulnerability discovery, and reverse engineering
- Modeling human behavior for cybersecurity and social simulation
- Theory and practice of distributed computing

**Artificial Intelligence**

- Social networks, analytics, and dynamics
- Bioinformatics
- Natural language processing
- Knowledge graphs
- Machine learning



### The ISI Workplace

ISI is committed to integrating diversity, equity, and inclusion into its operations and its community of students, faculty, and staff. The institute is dedicated to investing in tomorrow's leaders, providing internships, jobs, and program opportunities for students and unique opportunities for exceptional researchers from all backgrounds.

The DEI Steering Committee at ISI comprises senior researchers, the HR director, and USC DEI specialists. The committee hosted a focus group and followed up with a staff survey early in March. After reviewing the survey results, the institute made a significant financial commitment to its DEI efforts in advertising, staffing, recruiting, and inclusion training.

With guidance from an outside specialist, in 2022, ISI began a series of staff training sessions. Early sessions were open to all staff and featured interactive polls, videos, chat comments, and breakout sessions. Training for senior leadership followed in October and continues with monthly discussions. Likewise, open town hall meetings and discussion sessions are ongoing.

### ISI Outreach

In addition to supporting DEI in the workplace, ISI created a program to inspire underrepresented students to explore careers in science, technology, engineering, and math (STEM). 15 high school students and new grads participated in the free, four-week Stimulating STEM program, staying in USC housing and attending workshops, seminars, and lab tours at ISI.

The Stimulating STEM program was led by ISI's Jamani King and Lisa Bland, with support from a grant from Amazon. The participants saw various STEM research areas firsthand, including computer science, robotics, coding, artificial intelligence, cybersecurity, and space engineering. In many cases, they could apply their new knowledge to creating something functional, like building a robot out of Lego Technic bricks and connecting it to a battery to put it in motion.

The students received instruction from ISI researchers and faculty members, including space engineer David Barnhart, director of the Space Engineering Research Center (SERC). Under his guidance, they engaged with ISI grad students, faculty, and technology. When he described the space gripper "that looks like an octopus," all the students took note. The SERC director also showed a replica of the satellite he built with his team and sent into space in January.

One of the most pivotal effects of Stimulating STEM was the increased confidence in themselves. When asked how this program has helped define her future goals, one student stated: "Before, I was just thinking of getting my bachelor's, but the program has made me want to pursue a master's degree."







## Stephen L. Casner

Steve Casner passed away on July 4, 2022.

Steve joined ISI in the fall of 1973 after earning his BA in Mathematics from Occidental College. He completed his MS in Computer Science at USC in 1976.

Steve spent two decades at ISI, where he designed and implemented protocols and software for early experiments with packet voice using the ARPANET, the Advanced Research Projects Agency Network, predecessor to the Internet. He contributed to the first specification and implementation of packet voice on the Internet, the Network Voice Protocol. Steve and collaborators later developed the Real-time Transport Protocol (RTP) for delivering audio and video over IP networks in the 1990s, which drew significantly from the Network Voice Protocol and Packet Video Protocol developed at ISI. In 1992, he played an important role in establishing the worldwide Internet Multicast Backbone (MBONE). Steve and his colleague Eve Schooler

were awarded the IEEE Internet Award in 2021.

For a tribute to Steve and his many contributions, watch the remembrance at the 2022 IETF plenary session.



## Ronald B. Ohlander

Ron Ohlander passed away on January 10, 2022.

Ron joined ISI in 1985 and served as deputy director until he retired in 2006. Before joining ISI, Ron was a DARPA program manager in the Information Processing Techniques Office (IPTO), where he was deeply involved with funding artificial intelligence research. At ISI, Ron was critical in helping researchers obtain funding for their work.

ISI director Craig Knoblock noted—with admiration—  
”During his time in the service, the Navy sent Ron to Carnegie Mellon to do his PhD. He graduated from CMU in 1975 with a dissertation in artificial intelligence titled Analysis of Natural Scenes, which he completed in just over four years (a very short time for a CMU PhD).”

In 2004, Ron was honored with the Viterbi School of Engineering Service Award. Then dean of Viterbi, C. L. Max Nikias, remarked, “He has contributed to huge gains in research revenue for ISI; he has instituted many cost-cutting initiatives while also improving job security for our most talented individuals. Ron has also helped maintain excellent relations with DARPA and helped spin off the Internet Assigned Numbers Authority. Before coming to USC, he was an officer in the US Navy for two decades. He headed the basic AI research program at DARPA and received numerous medals and ribbons for his contributions.”



C. L. Max Nikias presents the Viterbi Service Award to Ron Ohlander (right)

## Artificial Intelligence Division

**YIGAL ARENS** INTERIM DIRECTOR | **ELIZABETH BOSCHEE** ASSOCIATE DIVISION DIRECTOR



The Artificial Intelligence Division is one of the world’s largest AI research groups, with over 170 faculty, research staff, and graduate students. Most researchers hold graduate degrees in computer science or related disciplines; many hold research faculty appointments at the Viterbi School of Engineering. The division conducts fundamental and applied research in these areas:

**Machine learning and data science** - developing efficient algorithms to analyze data from various application areas, including biomedical sciences, computational social science, and cybersecurity.

**Natural language processing** - low resource machine translation, multilingual representation learning, transfer learning, dialogue, decision-making, question answering, summarization, ontologies, information retrieval, extracting critical information about world events from foreign language sources.

**Knowledge graphs** - using AI and machine learning techniques to construct and exploit large-scale knowledge bases and to induce taxonomies from data.

**Scientific data analysis and discovery** - interactive knowledge capture, intelligent user interfaces, semantic workflows, provenance, large-scale data integration, analysis of biomedical data.



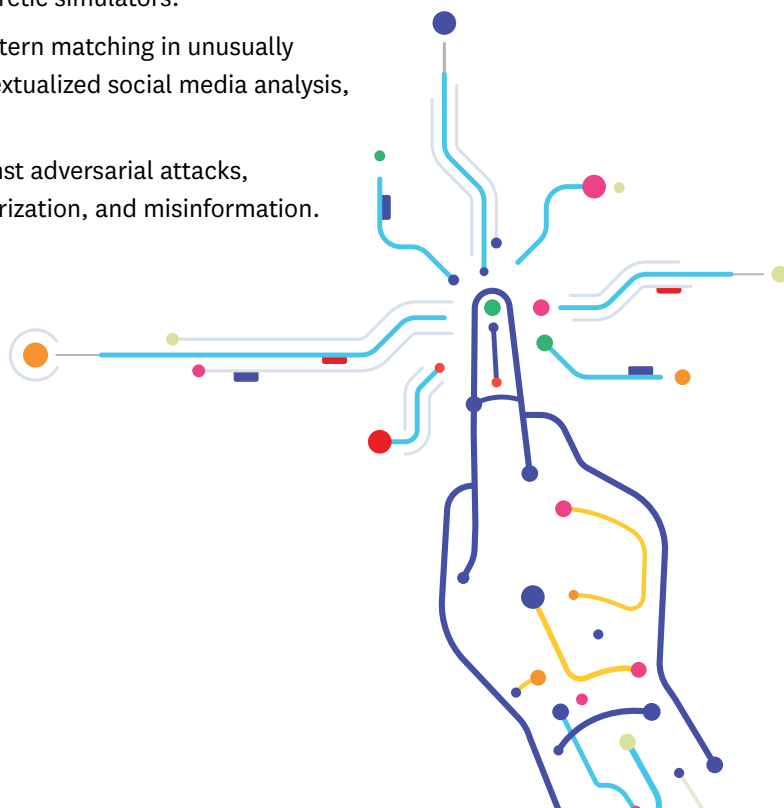
**Multi-modal understanding** - image and video understanding for deepfake detection, visual misinformation identification and multimedia analysis, face recognition, biometric anti-spoofing, learning “common sense” knowledge from text, video, images, and audio,

learning to automatically discover the structure of complex events from multilingual multimedia data.

**Common sense representation and reasoning** - building cognitively inspired computational paradigms for evaluating common sense AI, human-centric dialog agents, game-theoretic simulators.

**Computational social science** - structure detection and pattern matching in unusually complex systems with hidden information, large-scale contextualized social media analysis, detecting and responding to social engineering attacks.

**AI fairness** - detecting and mitigating bias, robustness against adversarial attacks, forecasting, crowdsourcing, identifying cultural values, polarization, and misinformation.



## Computational Systems and Technology Division

**STEPHEN CRAGO** DIVISION DIRECTOR



The Computational Systems and Technology Division (CS&T) focuses on:

- Heterogeneous cloud and embedded computing
  - Heterogeneous integration of novel materials, devices, circuits, and architectures for advanced microelectronics, photonics, and magnetic integrated circuits; hardware architectures for artificial intelligence and sensors
  - Numerical simulations of physical systems with high performance computing
  - Quantum computing, communication, sensing theory, and hardware
  - Reconfigurable computing and wireless networks
  - Science automation technologies
- Spaceborne and ground-based data processing
  - System software, including operating systems, runtime systems and frameworks, and compilers
  - Trusted and secure electronics and computing

The division’s researchers, research programmers, and graduate students represent disciplines including electrical engineering, computer science, physics, and math.

Ongoing initiatives include theoretical adiabatic quantum computing through the USC-Lockheed Martin Quantum Computing Center (QCC) and hardware security through ISI’s Secure and Robust Electronics Center (SURE). CS&T projects include system software for heterogeneous clouds and hardware-software design for novel chips and field programmable gate arrays. Division researchers explore applications and algorithms for processing large-scale and real-time streaming data and solving optimization problems; researchers work on system engineering for space systems.

CS&T teams are creating wireless networking and edge and fog computing technologies for battlefields and other challenging environments, along with social media platforms for people who lack trustworthy Internet access. Ongoing research on scientific automation tools enables scientists to focus on conducting science—instead of managing data. CS&T automation tools are used by astronomers, physicists, and earthquake specialists.

In the Laboratory for Quantum-Limited Information (QLIlab) at ISI Boston, CS&T researchers are exploring how to build the most sensitive communications and sensing devices when a signal is imprinted on only a fraction of a photon—which is the fundamental quantum particle of light.

At the Application Specific Intelligent Computing Lab (ASIC) researchers apply “alternate state variables” such as electrons, photons, phonons, and magnetic spins to deliver next-generation hardware fabrics using novel monolithic and heterogeneous device integration for artificial intelligence, smart sensors, smart and secure manufacturing, and quantum information. ASIC lab researchers and students bring hardware expertise in materials, devices, circuit design, and fabrication.

The Science Automation Technologies group explores the interplay between automation and the management of scientific workflows including resource provisioning and data management. The group leads the design and development of the Pegasus Workflow Management software and conducts research in job scheduling and resource provisioning in distributed systems, workflow performance modeling, provenance capture, and the use of edge and cloud platforms for science.

## Informatics Systems Research Division

CARL KESSELMAN DIVISION DIRECTOR



The research agenda of ISI’s Informatics Systems Research Division focuses on creating new types of sociotechnical systems that enable and accelerate discovery in domains of high societal impact. The Informatics Systems Research Division takes a holistic, systems-oriented approach, working in areas ranging from basic network service architectures, data management abstractions, computer security, and user-interface design to domain-specific algorithms. The division specializes in highly collaborative, user-driven research in the context of high-impact domain science.

In earlier work, the Informatics Systems Research Division developed grid computing infrastructures to support the

creation and operation of “virtual organizations” as a foundation for scientific collaboration and discovery. This work, which focused on understanding methods for sharing computing and storage, played a role in three Nobel prizes:

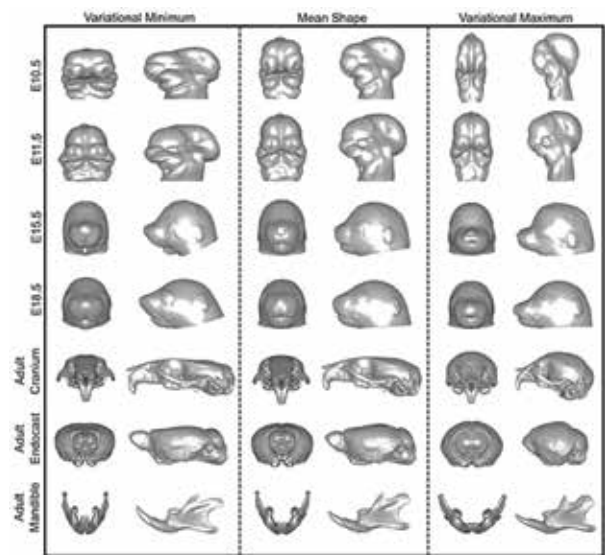
The Earth System Grid Federation contributed to the large-scale climate models for the International Panel on Climate Change. (2007 Nobel Peace Prize)

The data analysis for discovering the Higgs boson was performed on a global grid infrastructure, the Worldwide LHC Computing Grid. (2013 Nobel Prize in Physics)


The discovery of gravitational waves took place using the LIGO Data Grid and the Open Science Grid. (2017 Nobel Prize in Physics)

Researchers work closely with ISI’s artificial intelligence, networking, and distributed systems experts and with two of USC’s nationally ranked Viterbi School of Engineering departments: Computer Science and the Daniel J. Epstein Department of Industrial and Systems Engineering.

The division participates in collaborative projects with faculty in Dornsife College, Keck School of Medicine, and the Herman Ostrow School of Dentistry. Current collaborations range from basic science to clinical use cases in molecular biology, basic neuroscience, neuroimaging, stem cell research, and craniofacial dysmorphia.



Standardized mouse morphology data spanning numerous genotypes and developmental stages



Annotated mouse kidney image

## Networking and Cybersecurity Division

**TERRY BENZEL** DIVISION DIRECTOR | **DAVID BALENSON** ASSOCIATE DIVISION DIRECTOR



Networking is the backbone of the interconnected world, and cybersecurity is its guardian. The 70+ researchers, faculty, PhD students, and student workers in the Networking and Cybersecurity Division focus on cyber experimentation, network measurement, security for DoD networks, modeling human behavior, natural language understanding, autonomy in cyber-physical systems, binary program analysis, and distributed computing. Areas of research include:

**Cyber experimentation research, methods, and infrastructure** - creating models, frameworks, testbeds, tools, and approaches to enhance the science of cyber experimentation.

**Network measurement, analysis, and defenses** - developing novel Internet measurement techniques, tools, and data; applying formal methods and temporal logic for operational identification of attack dependencies and descriptions of traffic behaviors and machine learning for the classification and quantification of privacy leaks.

**Security for DoD networks** - improving the security of DoD networks through research into novel network embedding techniques, constraints validation, and protection of information in transit and from direct and side-channel attacks; applying machine learning to assist identification of applications based on their network traffic; applying advanced network tomography capabilities to identify the current state and structure of the network between communicating entities.



**Modeling human behavior for cybersecurity** - developing methods for detecting and fingerprinting attack campaigns by leveraging metadata from communication channels and

employing novel techniques to redirect attackers; the division's DASH agent simulation platform combines observational data with cognitive models of behavior and scales to millions of individual agents.

**Understanding communication for human-focused agents** - spanning text, visual, and structured modalities to use operator feedback for learning to interpret multimodal inputs and assist the operator; developing techniques to understand the language of cyberattacks and human interpretable representations of complex real-world events.

**Cyber-physical systems and autonomy** - studying the semantics of cyber-physical systems (e.g., space, energy, medical, automotive) to understand how software will interact with higher-level, human-understandable physical models, operate in mission-critical zero-trust environments, and interoperate with legacy systems.

**Binary program analysis for vulnerability discovery** - leveraging machine learning and reasoning to build state-of-the-art vulnerability discovery models that generalize vulnerability detection without manual labeling or extensive human effort.

**Theory and practice of distributed computing** - research on consistency and safety semantics, models of communication interaction among heterogeneous participants, privacy-preserving blockchain protocols for cryptocurrencies, energy-efficient and secure protocols, and robust and distributed network provenance for wide-area networks.

## The MOSIS Service

LIFU CHANG DIRECTOR



The MOSIS Service offers semiconductor fabrication services to semiconductor integrated circuit (IC) designers at universities, research organizations, defense and aerospace companies, and commercial design companies for multiple-project wafer and low-volume dedicated wafer projects.

Since its beginning in 1981, MOSIS (“Metal Oxide Semiconductor Implementation Service”) has processed more than 60,000 IC designs, averaging five chips a week.

MOSIS enables IC designers to prototype innovative semiconductor designs within an effective cost structure, offering a range of processes: CMOS FinFET, FD-SOI, Bulk, SiGe, high-voltage BCD, and other specialty processes. MOSIS collaborates with four major foundries—Intel Foundry Services, Samsung Foundry, TSMC, and GlobalFoundries—to offer a range of semiconductor processes from 12nm FinFET to 350nm.

MOSIS supports the design enablement, design support, and fabrication sign-off for Intel Foundry Services’ 22FFL and 16nm processes for digital, analog, and mixed-signal design flows utilizing Intel’s Process Design Kit (PDK), IP, and EDA tech files. MOSIS collaborates with Intel to establish a customer design support model connecting support capabilities from MOSIS, EDA companies, and Intel Foundry Services. In 2022, MOSIS provided support to nine riders on two tapeouts to Intel 16 foundry process under a major project sponsored by the Air Force Research Laboratory.

MOSIS is working with Samsung Foundry to promote joint projects in the Samsung Advanced Foundry Ecosystem (SAFE) Forum and other activities. MOSIS continues projects with TSMC and GlobalFoundries and is exploring new capabilities to add value to its offerings for Extreme UltraViolet (EUV) processes from TSMC.

In 2022, the USC Ming Hsieh Department of Electrical Computer Engineering, MOSIS, and the Industrial Technology Research Institute (ITRI) in Taiwan announced a new collaboration to advance the field of microelectronics in research and development, IC design, and manufacturing. MOSIS and ITRI will concentrate on IC design and development service, intellectual property support, and silicon shuttle service.

Also this year, MOSIS and WIN Semiconductors signed a memorandum of understanding to accelerate the development of GaAs and GaN Monolithic Microwave Integrated Circuits (MMICs) with WIN Semiconductors’ GaAs and GaN process technologies. Project flows for customers to utilize MOSIS to fabricate MMICs at WIN Semiconductors are now being released.

In 2022, MOSIS and SkyWater Technology began a collaboration to enable design-to-manufacturing for IC prototyping and production utilizing SkyWater’s CMOS process technologies. MOSIS will provide multi-project wafer (MPW) shuttle access to SkyWater’s S90 and S130 mixed-signal CMOS technologies. MOSIS and SkyWater teams are preparing for the MPW workflows.

In addition to the commercial service, MOSIS staff conduct research in silicon fabrication and Design for Manufacturability (DFM) areas, producing academic papers and patents.



THE MOSIS SERVICE



Watch a quick introduction to the MOSIS Service.

## Analysis of Network Traffic: The ANT Lab

JOHN HEIDEMANN FOUNDER




The goal of the Analysis of Network Traffic Lab (ANT Lab) is to improve the Internet by discovering new ways to understand network topology, traffic, security, use, and abuse.

ANT was originally founded in 2000. The group’s research is led by multiple principal investigators in ISI’s Networking and Cybersecurity Division in collaboration with principal investigators from the University of Memphis and Columbia University.

The lab carries out Internet measurement with active and passive data collection and big data analysis of network data.

ANT has received support from the Department of Homeland Security, the National Science Foundation, DARPA, and from Cisco, Verisign, Northrup Grumman, Amazon, and a Michael Keston Exploratory Research Award.



Fall 2022 Advanced Computer Communications networking class poster session, sponsored by the ISI ANT Lab

## Center on Artificial Intelligence Research for Health

MICHAEL PAZANNI DIRECTOR

WAEEL ABDALMAGEED CO-DIRECTOR

ABIGAIL HORN CO-DIRECTOR

JOSE LUIS AMBITE CO-DIRECTOR

CARL KESSELMAN CO-DIRECTOR



**AI4HEALTH WORKSHOP**  
BY USC ISI • NORRIS CANCER CENTER  
**DECEMBER 1**  
AT 11:00AM

LINKING RESEARCHERS ACROSS HEALTH AND AI

We invite you to the first joint workshop on AI and health, co-hosted by the Center for AI Research for Health, USC ISI, and the USC Norris Cancer Center. The workshop is aimed at identifying and pursuing potential avenues at USC for health applications of artificial intelligence. It will be an opportunity for health and AI scientists to connect and brainstorm how AI methods can be used to open up new avenues to critical, challenging health problems.

**AI RESEARCH THEMES**

- MACHINE LEARNING
- EXPLAINABLE AI
- IMAGE CLASSIFICATION
- HEALTH INFORMATICS
- AUTOMATED DATA SCI
- FEDERATED LEARNING
- WEARABLE SENSORS
- NEUROIMAGING
- MEINFORMATION
- GENE EXPRESSION

**HEALTH DOMAINS**

- MULTI-OMICS DATA
- ENVIRONMENTAL
- HEALTH
- GENOMICS
- SPATIAL EPIDEMIOLOGY
- CLINICAL TRIALS IN
- CANCER
- FOOD SYSTEMS

**PRACTICAL INFORMATION**

WHO: Open to the USC community - register at [ai4health.usci.edu](https://ai4health.usci.edu)  
 WHEN: Thursday December 1, 11:00am-1:00pm  
 WHERE: Health Sciences Campus, Anschutz Conference Center, Room LC 550, Norris Cancer Research Tower (Building NRT), 1441 Eastlake Ave., Los Angeles CA 90033  
 CONTACT: Abigail Horn [ahorn@usc.edu](mailto:ahorn@usc.edu), David Crand [dcrand@usc.edu](mailto:dcrand@usc.edu)

The center on AI Research for Health was created to nurture collaborations between researchers in AI and those in the health sciences. Focus areas include:

- Data management
- Precision health
- Machine learning for health
- Telehealth
- Knowledge discovery and data analytics

The center sponsored the First Joint Workshop on AI Research for Health, co-hosted by the USC Norris Comprehensive Cancer Center, in December to identify and pursue potential avenues at USC for leveraging innovations in AI to impact health and make new discoveries in critical, challenging health problems.

## Center for Cyber Defense Technology Experimental Research

TERRY BENZEL CO-DIRECTOR

JELENA MIRKOVIC CO-DIRECTOR



The center focuses on cyber experimentation research, methods, and infrastructure—catalyzing a robust ecosystem of experimental infrastructure, capabilities, and communities—for research. Center researchers perform research into models, frameworks, testbeds, tools, and approaches to enhance the science of cyber experimentation and build experiments that are rigorous, reusable, and repeatable.

The center operates DeterLab and associated research infrastructure. The DeterLab Cybersecurity Experimentation Testbed is a public resource serving over 1,000 research teams from 46 countries. Additionally, the DeterLab has been extensively used for cybersecurity education and impacted more than 20,000 students. Numerous publications, master’s projects, and PhD theses have resulted from the use of DeterLab, which has been operating since 2004. The Deter testbed now offers three custom testbeds: DCOMP, a specialized distributed computing testbed; Searchlight Testbed, for enterprise control of application QoS at scale; and a Modernized DeterLab.

## Center on Knowledge Graphs

JAY PUJARA DIRECTOR



The Center on Knowledge Graphs creates new approaches for amplifying artificial intelligence using structured knowledge. The group combines expertise in artificial intelligence, machine learning, the Semantic Web, natural language processing, databases, information retrieval, geospatial analysis, business, social sciences, and data science. The center comprises 13 senior ISI researchers, guiding the work of 29 PhD students, 50 master’s and undergrad students, and three research engineers.

The center has built tools and knowledge graphs to address challenging, real-world problems such as reducing global food insecurity, fighting human trafficking, assessing medical and clinical data, fostering pharmacological discovery, and ensuring scientific reproducibility.

## CI Compass: An NSF Cyberinfrastructure Center of Excellence for Navigating the Major Facilities Data Lifecycle

EWA DEELMAN DIRECTOR



CI Compass provides expertise and active support to cyberinfrastructure practitioners at NSF Major Facilities, such as the Laser Gravitational Wave Observatory (LIGO), the National Ecological Observatory Network (NEON), and the Regional Class Research Vessel (RCRV). These facilities develop and operate unique scientific instruments, data, and computational capabilities for scientists, engineers, students, and the public.

CI Compass provides advice, design choices, and proof of concepts related to the Major Facilities Data Lifecycle. This multi-university collaboration, which includes researchers from Indiana University, RENCI, the University of Notre Dame, Texas Tech, and the University of Utah, contributes knowledge and expertise to enhance the NSF cyberinfrastructure ecosystem.



## Laboratory for Quantum-Limited Information


**JONATHAN HABIF** DIRECTOR



The Laboratory for Quantum-Limited Information (QLIlab) is dedicated to understanding and demonstrating the fundamental physical limits for extracting information from physical signals using quantum mechanics. The work at the lab focuses on developing revolutionary designs for information processing systems to enable communication and sensing at the fundamental limits of nature.

A complete understanding of a physical signal (such as light or radio frequency) is possible only when quantum mechanics is employed to mathematically describe the signal. The quantum mechanical description provides insight into the maximum amount of information that could possibly be extracted from the physical signal. Lab researchers calculate these fundamental limits for performing information processing tasks and build laboratory experiments to demonstrate the ability to achieve these fundamental limits.

The QLIlab is located at ISI's Boston location and offers a multi-disciplinary training ground for students interested in physics, mathematics, and engineering.



## USC-Lockheed Martin Quantum Computing Center

**DANIAL LIDAR (USC)** SCIENTIFIC & TECHNICAL DIRECTOR

**FEDERICO SPEDALIERI** OPERATIONAL DIRECTOR



A joint effort of Lockheed Martin Corporation and USC, the Quantum Computing Center (QCC) is housed at ISI. Faculty, researchers, and students perform basic and applied research into noisy, intermediate-scale quantum (NISQ) computing devices and collaborate with researchers around the world.

The Quantum Computing Center was the first organization outside D-Wave Systems to house and operate its own D-Wave quantum computer, and it has conducted pioneering research on four different generations of these early NISQ processors.

Quantum computing potentially offers orders-of-magnitude gains in speed and memory, as well as greater security, for some computing problems. Applications include “big data” analysis, classification, optimization, and machine learning.

## Science Automation Technologies: SciTech

EWA DEELMAN DIRECTOR



The Science Automation Technologies group (SciTech) conducts research and provides technologies that empower the scientific community to efficiently conduct complex computations on the national, campus, and industrial cyberinfrastructure. By raising the level of abstraction for computation specification (abstract scientific workflow) and by designing workflow management systems, scientists are able to focus on their research questions rather than the details of the cyberinfrastructure. The group develops open-source tools that provide the computational foundations enabling scientists to seamlessly run their experiments and analyses in local and distributed resources.

The group concentrates on these research areas:

- Workflow management systems for science
- Machine learning for science
- Enhancing national cyberinfrastructure
- Modeling and simulation of distributed computing systems

---

## Space Engineering Research Center

DAVID BARNHART DIRECTOR



The Space Engineering Research Center (SERC) is dedicated to disruptive space engineering, research, and education for the second-generation workforce—including hands-on training to build, test, and fly spacecraft and satellites. SERC hosts high school, community college, undergrad, graduate, and PhD students from the US and worldwide.

SERC is a longstanding joint venture between ISI and the USC Department of Astronautical Engineering.

SERC seeks to challenge traditional methods of space R&D, manufacturing, and exploration with approaches that dramatically reduce costs, enable novel capabilities, and support the creation of space-to-space services and interactions.



## STEEL: Security Research Lab

**JELENA MIRKOVIC** DIRECTOR

**LUIS GARCIA** CO-DIRECTOR

**GENEVIEVE BARTLETT** CO-DIRECTOR

**CHRISTOPHE HAUSER** CO-DIRECTOR



Founded in 2012, STEEL researchers, PhD students, and master's students conduct cutting-edge research in cybersecurity and testbed experimentation. The lab focuses on research in network-based cyber attacks, human aspects of cybersecurity, embedded cybersecurity, and binary analysis. STEEL also contributes to improving security through cyber testbed experimentation, including work in making experimentation more rigorous, more accessible, and reliably repeatable.

## Secure and Robust Electronics Center

**MATHEW FRENCH** DIRECTOR



The Secure and Robust Electronics Center (SURE) focuses on secure, robust hardware R&D to help protect national security, advance industry capabilities and competitiveness, and ensure that citizens can rely on the integrity of their devices. SURE researchers perform applied research to make chips trustworthy, secure, resilient, and reliable.

The center investigates state-of-the-art manufacturing complexities that have compromised critical aspects of chip production. The more powerful integrated circuits become, the more opportunities arise for their integrity to be compromised. SURE researchers also aim to accelerate the scale, pace, and impact of hardware robustness and technology development.

The center's work has been supported by government agencies and by industry partners, including Xilinx Inc, Altera Corporation, Synopsys, and Cadence. Academic collaborators include Stanford, Virginia Tech, Brigham Young University, Arizona State, University of North Carolina, and Georgia Tech Research Institute.

## Visual Intelligence and Multimedia Analytics Laboratory

**WAEEL ABDALMAGEED** FOUNDING DIRECTOR



In the Visual Intelligence and Multimedia Analytics Laboratory (VIMAL), researchers at all three ISI locations work on visual misinformation identification, face recognition, biometrics, and multimedia forensics.



VISUAL INTELLIGENCE AND MULTIMEDIA ANALYTICS LABORATORY

VIMAL's research addresses challenging problems in security, surveillance, and multimedia content analysis, including improving the security of biometric systems, making deep learning algorithms more robust against adversarial attacks, and identifying deepfake videos and manipulated images.

## KESTON EXPLORATORY RESEARCH AWARDS

ISI received a generous endowment gift in November 2015 from Los Angeles entrepreneur, philanthropist, and engineer-at-heart Michael Keston and his wife and philanthropic partner, Linda Keston. The Kestons endowed ISI's Keston Executive Director chair.

In addition to endowing the director's chair, a portion of the income from the Keston endowment is dedicated each year to sponsoring the Keston Exploratory Research Awards, intended to foster and support exploratory early-stage research not yet funded by outside sponsors. ISI has contributed additional funds to expand the program, creating the ISI Exploratory Research Awards.

The Keston and ISI Exploratory Research Awards support research projects that are intellectually intriguing, significant to society, and have the potential to produce results within a year or so.

The following describes the work carried out in 2022 under the aegis of the Keston and ISI Exploratory Research Awards program.



### *IRIS: Integrated Retinal Functionality in Image Sensors*

**Akhilesh Jaiswal, Ajey Jacob, Gregory Schwartz** (Northwestern University),  
**and Maryam Parsa** (George Mason University)

Today's computer vision exclusively relies on light intensity-based pixel data collected through CMOS image sensors. However, the appropriate context for the pixels is missing (or vague) with respect to real-world events. The *biological* retina, in contrast, provides dozens of feature-selective computations such as object motion, estimated trajectory of objects, object velocity, and shapes. The IRIS project works on embedding retina-like computations into image-sensing camera technology. The team of engineering researchers, retinal neuroscientists, and ophthalmology specialists modeled the computational processing in downstream retinal circuits and leveraged potential 3D stacking of semiconductor chips to chalk out a pathway toward developing new retina-inspired sensors.

### *Unbiased and Explainable Diagnosis of Melanoma*

**Michael Pazzani and Mohammad Rostami**

Malignant melanoma is the fifth most common form of cancer diagnosed in the US. Despite being more prevalent among Caucasians, it has a higher mortality rate in people of color because it is diagnosed at later stages. The researchers developed an AI-based approach for automatically differentiating melanoma moles from benign moles at early stages using images taken by ordinary smartphones. The method is explainable because it relies on extracting and localizing clinical biomarkers on the input images to diagnose melanoma. As a result, clinicians can interpret the results, which is crucial to adopting AI in healthcare. To mitigate bias, synthetic images that represent minority sub-populations are generated to offer a more balanced dataset than existing datasets.

### *Electronically Programming Thermal Emissions and Signatures from Warm Bodies*

**Jonathan Habif**

Researchers at the Laboratory for Quantum-Limited Information (QLIlab) developed a revolutionary electro-thermo-optic device able to electronically program and switch the rate at which Planck radiation is emitted by warm bodies—enabling the ability to control heat expulsion from the surface of an object. This early-stage technology development could apply to many areas, including a new communications signaling medium or a technique for engineering heat flow across the surface of an object. The team built a prototype unit that can be shared with engineers or entrepreneurs so that they can demonstrate new technologies that capitalize on this new capability. As an exemplar demonstration, they used the radiators to demonstrate a novel optical communications technique achieving 100 bps data rates over laboratory-scale distances.

### *Standing on the Shoulders of Giants: Understanding Creativity and Collaboration through Temporal Knowledge*

*Graph Learning*

**Jay Pujara**

How are great ideas born? Do they spring from individual geniuses, large teams, or multidisciplinary collaborations? Although models have been proposed in studies of innovation, analyses have been limited in the size and scope of the data. The ISI team, which included four graduate students, studied models of collaboration with hundreds of millions of research publications to identify patterns of innovation. Formulating publication data in a massive science temporal knowledge graph of authors, publications, institutions, and journals, they identified scientific communities and how they evolved over time. They also built forecasting models for new research ideas. Applying a causal model, they found that authors who diversified their collaborations and communities were more successful. The project created resources for exploring the lineage of ideas, collaborative structures, and the evolution of individuals' careers.

### *Why Resilience in Innovation Is Necessary and How to Foster It*

**Kristina Lerman, Jay Pujara, and Keith Burghardt**

Resilience helps people and organizations recover from shocks. The COVID-19 pandemic offered a unique natural experiment to study resilience because the impact was highly uneven, with policies varying by country, state, community, and even institution. The team studied the heterogeneous impact of the pandemic on scientific innovation by analyzing newly available bibliographic data sets covering millions of authors and papers. They applied methods from the field of causality to understand how the disruptions affected research collaborations and the factors that promote resilient research. They learned that research activity appears to have accelerated, but collaborations decelerated. There was a decline in the number of active junior authors yet an acceleration in active senior authors. The research showed how the impact of the pandemic on researchers varied by field, career stage, and gender.

*ISI researchers make fundamental contributions to basic and applied computing research and engineering. The research carried out at ISI spans many topics, including:*

- Machine learning
- Quantum physics
- Small satellites
- Natural language processing
- Chip and hardware security
- Bias in job ads, literature, and social media
- Network security

*The research highlights that follow present a sampling of ISI's work.*

### Reproducible Video Teleconferencing Traffic Generation

Large-scale video teleconference (VTC) environments are commonplace today as more people work from home and participate in meetings and conferences virtually. However, tools to model human-like VTC traffic for systematic and disciplined simulation and emulation are lacking.


Researchers in the Networking and Cybersecurity Division built a VTC traffic generation tool that enables experimenters to create representative traffic on-the-wire with human-like bot interaction, including “talking” and “listening” periods. These interaction patterns are extended to create representative VTC scenario templates, including small and large team meetings, workshop discussions, conferences, and keynote-speaker setups with multiple receivers and multiple mixed modes for interaction, including chat, audio, and video.

The VTC tool is completely automated and reproducible and can be deployed in any testbed environment to benchmark emerging networking technology, such as network traffic classification and quality of service (QoS) management. The VTC traffic generation tool can also be coupled with client applications from commercial VTC solutions, such as Zoom and Teams, to integrate human and bot-based interactions.

The VTC traffic generation tool was developed for the DARPA Searchlight program for managing QoS for distributed applications on the Internet. The source code is open source and available for download.



**Contact: Alefiya Hussain, [hussain@isi.edu](mailto:hussain@isi.edu)**



*Realistic video teleconference (VTC) traffic generated with eight bot clients attending a meeting*

## Robust and Explainable Methods for Commonsense Reasoning in the Real World

People apply commonsense knowledge and reasoning to comprehend everyday situations. Stories, questions, and arguments about these situations often leave out assumed, or background, knowledge. If a child feels cold, we might suggest that the child should grab a blanket, implicitly using background knowledge that blankets increase a person’s warmth and that being cold is uncomfortable. If the child is, for example, camping and cannot find a blanket, we may suggest using a jacket instead, based on the similarity of jackets to blankets as warm, safe, and often nearby. Such commonsense reasoning is easy and intuitive for people, but it has long been considered “black matter” for machines. Commonsense reasoning is challenging, even to conceptualize.

Recent research on neural large language models, evaluation benchmarks, and crowdsourced knowledge sources has enabled many evaluation datasets to be solved by “black box” models that are adapted to representative training data. But, benchmark-specific training data is rarely available in realistic scenarios, and language models alone cannot generalize well to novel use cases. Moreover, in order to be accepted, machines that collaborate with people need to be transparent about their reasoning.

How can methods that exhibit commonsense reasoning on unseen tasks be developed and explained?

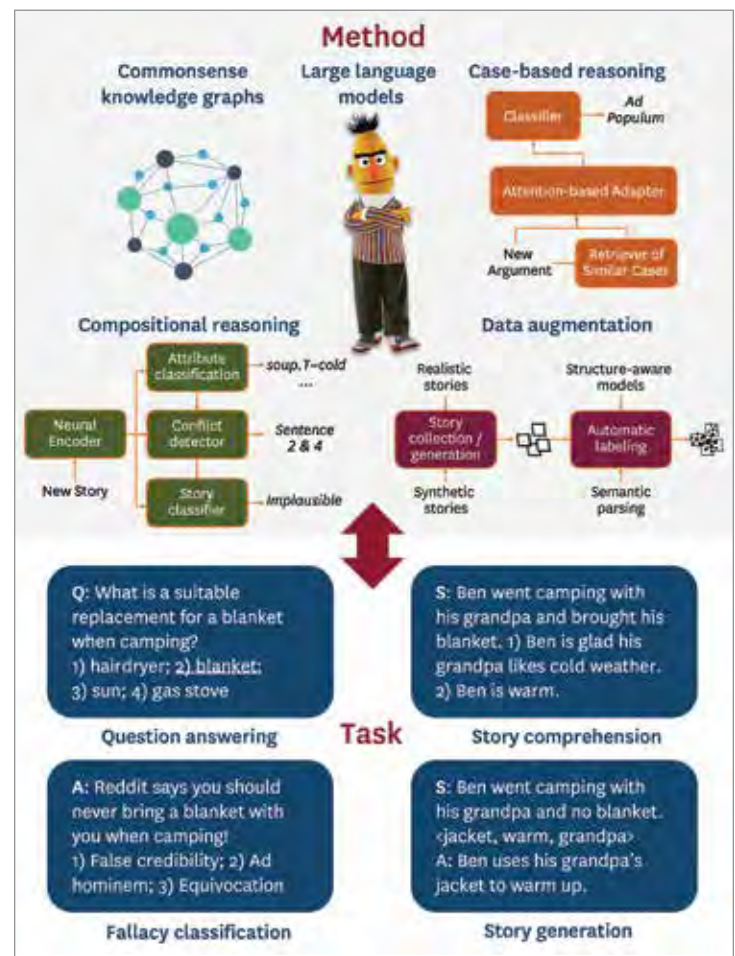
Researchers in ISI’s Artificial Intelligence Division developed a range of methods that perform robust and explainable reasoning on tasks designed to rely on commonsense reasoning, such as question answering, story understanding, story generation, and logical fallacy classification.

To answer questions on unseen benchmarks, the team developed a zero-shot method that generates synthetic data from a large, consolidated knowledge graph and uses the method to reason about novel questions. More recently, the researchers developed PINTO, a method that enhances reasoning by prompting knowledge from very large neural language models, like GPT-3, and uses counterfactual training for better robustness.

The researchers also implemented the classical method of case-based reasoning by using language models as case retrievers and adapters. They applied this novel case-based reasoning method to classify fallacious arguments and again leveraged prompting large language models to obtain additional information about counterarguments, goals, structure, and explanations of an argument. Moving from questions and simple fallacies to more complex structures, they developed methods that automatically process novel stories. Initial experiments showed that state-of-the-art models could not connect two stories by drawing analogies between their structures. They then developed CGLI, a method that extracts and leverages intermediate representations of stories, namely tracking of participant states (e.g., soup.temperature = cold) and conflict detection (e.g., sentence 3 is contradictory to sentence 5). A follow-up work, named LEAP, coupled CGLI’s explainability with the automatic generation of synthetic stories from knowledge graphs and an automatic labeling system that can annotate both synthetic and natural stories. The combination of these techniques led to high accuracy on novel story tasks, together with native explainability due to the compositional reasoning in LEAP and CGLI. They also developed a method that generates new stories by “imagining and verbalizing” (I&V) possible stories into scene graphs, which again relies on combining language models with graph structures through data augmentation and adaptations in the model architecture.

The research showed that combining neural language models with knowledge sources can lead to models with higher robustness and better explainability.

Contact: Filip Ilievski, [ilievski@isi.edu](mailto:ilievski@isi.edu)



A framework for robust and explainable methods designed to generalize to open-world tasks

## Unraveling New Materials with Novel Quantum Many-body Simulation Methodologies

One of the greatest challenges that physicists, chemists, and material scientists face is how to gain insights into the behavior of complex materials, such as strongly correlated electron systems or quantum spin glasses, in extreme conditions, for example, at extremely low temperatures or in the presence of high-intensity magnetic fields.

Recreating such complex materials in the lab is extremely costly. The best tool available is to *simulate* these materials on high performance computers. Although this is feasible for simple materials under normal conditions, it is exceedingly complicated when larger and more complex structures are involved because quantum phenomena emerge—behaviors that are not only inconsistent with our intuition but also obey the far more complex and counterintuitive laws of quantum mechanics. Capturing the salient properties of complex materials, such as phase transitions or response functions, requires tremendous computational powers that overwhelm even the largest supercomputers.

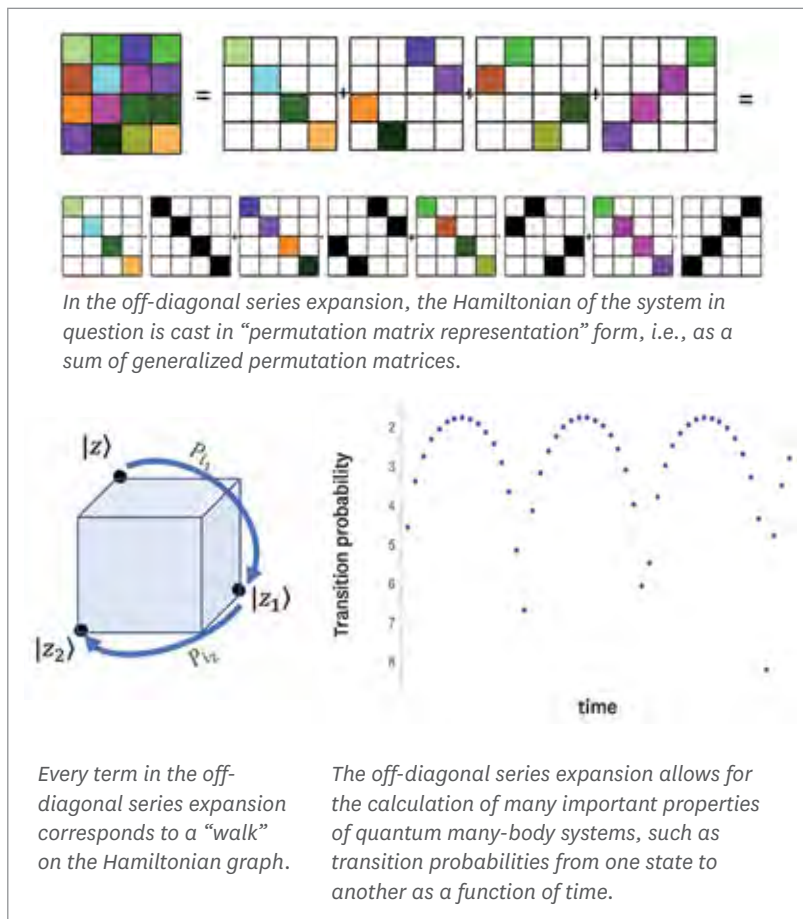
New algorithms are needed, algorithms that are more efficient and capable of capturing emergent quantum phenomena but that do not consume too much time or memory resources.

The ISI computational physics group has developed novel algorithms to efficiently simulate a broad range of systems, such as magnetic materials and strongly correlated fermions and bosons on a lattice (or mixtures thereof). The group devised, coded, and implemented new parallelizable algorithms for the efficient simulation of quantum many-body systems. The algorithms are based on mathematical advancements that introduce a novel power-series expansion for the matrix exponential and other matrix functions, which is termed “off-diagonal series expansion.”

The algorithms are considered state-of-the-art in the field of computational physics. The researchers have successfully employed them to provide a new understanding of several longstanding problems in physics, such as the infamous sign problem, quantum spin glasses, and certain “frustrated” systems—systems in which the constituent particles arrange themselves in non-trivial positions or configurations due to the existence of opposing forces.

Another arena where these techniques are applicable is quantum computers—at present, mostly theoretical devices that, unlike standard computers, operate in accordance with the laws of quantum mechanics, the laws that govern the behavior of the elusive hard-to-simulate complex materials. As such, quantum devices are thought to serve as a more natural platform for the simulation of quantum materials. Although fully functioning quantum devices do not yet exist, the scientific community is devising algorithms in preparation for their future existence. ISI’s computational physics group has successfully adjusted and ported its algorithms into the challenging realm of quantum computing. These quantum algorithms will allow scientists to more precisely simulate quantum materials on quantum computers when they eventually become ubiquitous.

**Contact:** Itay Hen, [itayhen@isi.edu](mailto:itayhen@isi.edu)

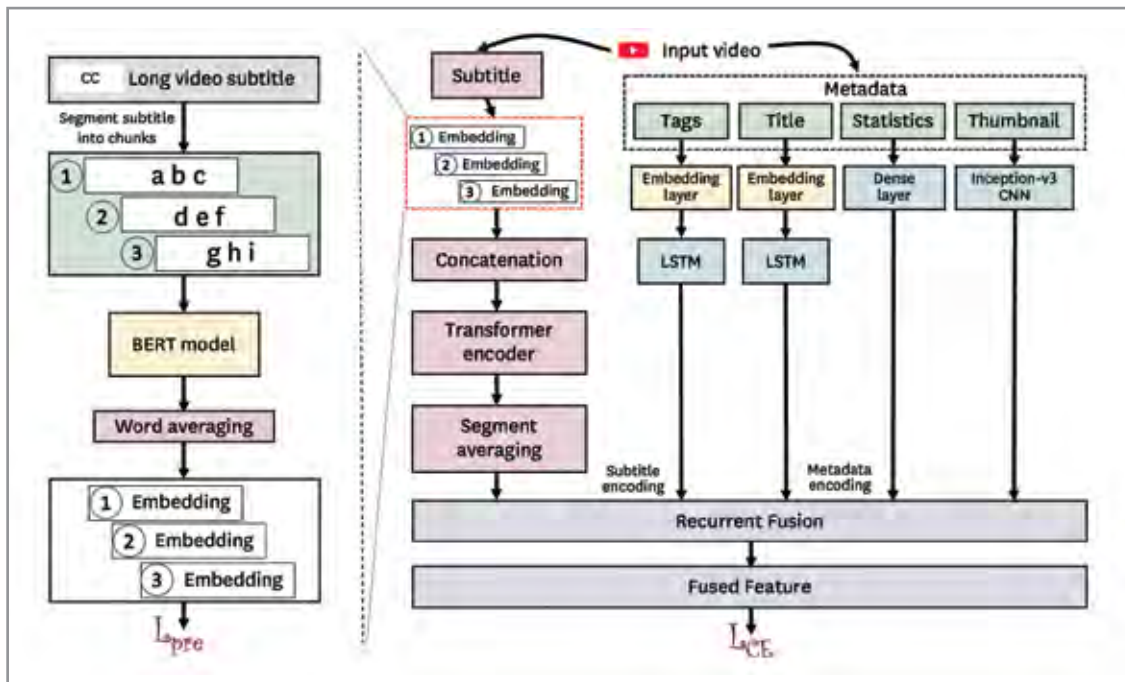




## Protecting Kids from Harmful YouTube Content

YouTube videos are a popular way of disseminating content, and they appeal to a diverse audience. Along with adults and older children, children as young as toddlers are avid consumers of YouTube videos. But young children often cannot evaluate if content is appropriate for their age, and parents have very limited options to enforce content restrictions on YouTube. Young children can thus become exposed to inappropriate content. This threat is exacerbated by the presence of YouTube poops—videos that mix appropriate content with inappropriate content inserted by the video creator.

ISI researchers and collaborators investigated how to build an automated classifier for YouTube videos that can flag videos inappropriate for young children. They collected a large, balanced dataset of 70,000 appropriate and inappropriate videos. (The dataset has been publicly released, so others can reproduce the results and further investigate video classification.) The team explored the use of video metadata (e.g., title, tags, thumbnail, likes, dislikes, views) and video subtitles in the classification process. They compared baseline machine learning approaches against novel approaches, including the BERT language model, and against their own approach, named Samba.



Training the Samba model

To train the Samba model, they fine-tuned BERT on segmented fixed-length chunks of a given subtitle. Next, they used contrastive learning on BERT segment embeddings to help their model learn similarities between subsequent segments in the same subtitle. The chunk embeddings produced by the fine-tuned BERT were fed into the Samba Transformer Encoder. The chunk encodings were averaged to obtain one fixed-length encoding for a given subtitle. Finally, a feature fusion was performed to combine the embeddings into a single embedding responsible for the classification.

Samba outperformed all metadata-based and subtitle-based approaches, achieving superior accuracy (95%), precision (94%), and recall (96%). This is because the model captures the context and semantic meaning of an entire subtitle by learning relationships between pairs of segments via contrastive learning and by learning how metadata and subtitle embedding relate to the final classification of the video.

The researchers envision the Samba classifier becoming part of a browser plugin, enabling parents to protect their children from harmful content on YouTube. They also believe that vulnerable populations, such as people contemplating suicide, people with eating disorders, or people with PTSD, could be protected from harmful YouTube content by applying Samba to a balanced dataset containing appropriate and inappropriate videos for that vulnerable population.

Contact: Jelena Mirkovic, [mirkovic@isi.edu](mailto:mirkovic@isi.edu)

## “I Know What You Did on Venmo” Quantifying the Dangers of Public-by-Default Platforms

Researchers in the Networking and Cybersecurity Division, working with University of Texas, Austin collaborators, measured the prevalence of sensitive content—and risk—in Venmo payment notes. These notes are public by default, and users tend to reveal sensitive information in them, such as their relationships, locations, sexual preferences, race, phone numbers, friends, and alcohol use.

The researchers completed the largest quantitative study of its kind detailing how millions of users reveal extremely personal information about themselves on Venmo. They examined 389,000,000 public notes over an eight-year period from 2012 to 2020.

Using a machine learning model, the researchers classified information contained in transaction notes as sensitive or non-sensitive. They further refined the data by grouping sensitive information into 14 categories, including criminal and violent behavior, sexual orientation, health, and physical location. They found that 41,000,000 transaction notes, or 10.5% of the public notes, leaked “some sensitive information such as [a] health condition, political orientation or drug and alcohol consumption.” Nearly 40% of the users had publicly shared sensitive information on the financial app at least once, in many cases inadvertently.

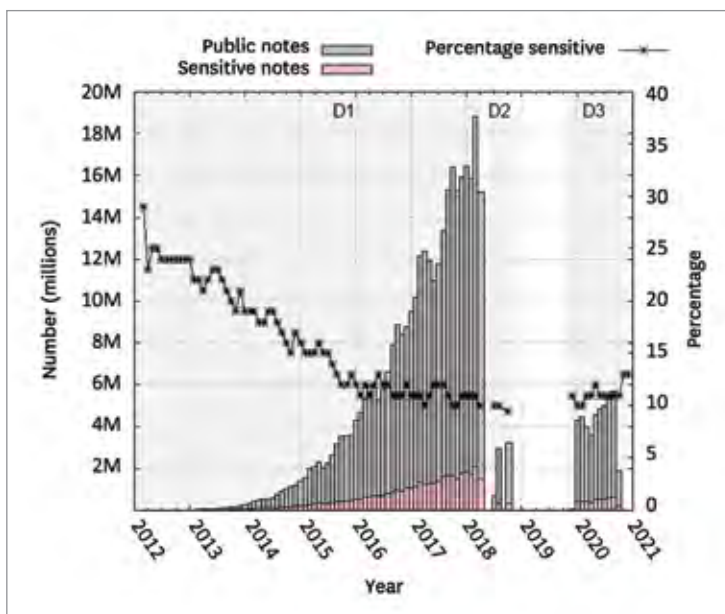
Although the percentage of users and transactions with public notes declined over time, the absolute number increased due to the large influx of new Venmo customers.

The researchers showed that it is possible to infer sensitive group membership, such as Alcoholics Anonymous, based on the sensitive note content of a small number of members. For example, if a group had an innocuous name, but a handful use AA terminology, then one could infer that it is an AA group, and everyone’s AA membership would be revealed.

During their research, the researchers attempted several times to speak to Venmo, which is owned by PayPal, but no high-ranking official responded. The team did receive money through PayPal’s Public Bug Bounty program for finding multiple security flaws in Venmo’s APIs. Venmo has since addressed the security flaws, but privacy issues remain.

The researchers advise Venmo users to make their profiles private, and to make all past and future Venmo notes private—with one click in their Venmo profile.

**Contact:** Jelena Mirkovic, [mirkovic@isi.edu](mailto:mirkovic@isi.edu)



Total public notes and sensitive notes over time

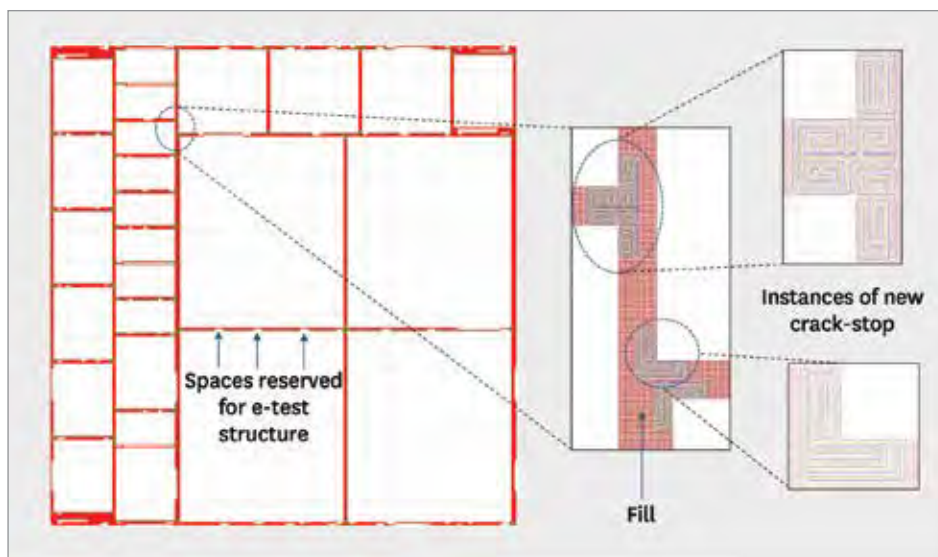
## Optimizing Yield in Three-Dimensional IC Fabrication through Insertion of Stress-Reducing Material in the Die Assembly Process of a Silicon Wafer

The MOSIS Service and Xyalis EDA Company collaborated in 2022 to produce a multiple project wafer (MPW) reticle placement flow that optimizes the yield for the dies going through three-dimensional IC (3DIC) fabrication, such as Through-Silicon Via. The flow protects the dies from fractures caused by thermal-mechanical stress in the wafer thinning process for 3DIC.

Polyimide material, which is flexible and can alleviate stresses, is carefully inserted in the spaces between dies in groups of pre-calibrated, stress-reducing-effective patterns. The algorithm for inserting the polyimide patterns and the mechanism of the patterns serve to minimize die fracture.

The MOSIS Service continues to perform research on 3DIC technologies that can be extended from this MPW reticle placement flow.

Contact: Lifu Chang, [lchang@mosis.com](mailto:lchang@mosis.com)



Example of inserting crack-stop (stress reduction) polyimide patterns in the die frame. Each white-color rectangle is a die. Serpentine patterns of polyimide are filled in the empty spaces between the dies.

## Understanding Gender Bias in Literature

ISI researchers utilized AI technologies to conclude that male characters are four times more prevalent in literature than female characters.

The researchers were motivated by current work on implicit gender biases. Although many published studies survey and analyze the qualitative aspects of female representation in literature and the media, the ISI researchers collected *quantitative* data employing machine learning algorithms. They accessed the Gutenberg Project corpus to analyze 3,000 books of English literature from 1800-1922, the pre-modern era, including novels, short stories, science fiction, and mysteries.

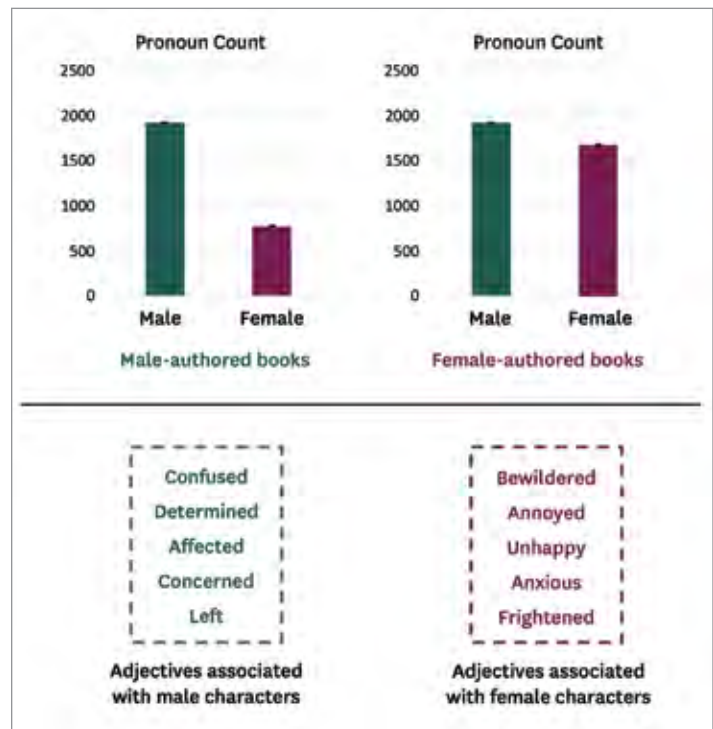
Their study outlined several methods for defining female prevalence in literature. They utilized Named Entity Recognition (NER), a natural language processing method used to automatically extract gender-specific characters using deep neural networks. Once the characters were extracted, they used machine learning to determine whether a character was male or female. To ensure that a single definition of prevalence does not bias the study, they also looked at alternative definitions of prevalence, e.g., the number of female pronouns in a book compared to male pronouns, as well as the number of books where the main character is female.

Using advanced statistical methods first developed in the computational social sciences, they found a 4:1 difference between male and female character prevalence regardless of how prevalence was defined (pronouns, named characters, or the main character). This difference was found to persist over their period of study. However, a positive finding from the study was that the discrepancy between male and female characters decreases by more than half under female authorship, showing that women represented their gender in their writing much more than male writers did. Rigorous statistical analysis showed that the difference was highly significant, providing strong evidence for having greater diversity in our cultural discourse.

The diversified methods to measure and determine female representation in literature are limited, however, when authors are neither male nor female. AI tools for identifying plural words, such as “they,” which may refer to a non-dichotomous individual, do not yet exist. Still, the study’s findings build the framework for approaching such social issues and building the technologies that can address these deficits.

The study provides a blueprint for future work on quantifying the qualitative findings they discovered. Without the inherent bias from human-designed surveys, the natural language processing technology enabled them to find adjective associations with gender-specific characters. Specifically, they used deep neural networks and unsupervised clustering techniques to automatically extract and group such associations based only on the context surrounding around each character’s occurrence in the text. Upon analyzing the groups, they found that the words associated with women were adjectives such as “weak,” “amiable,” “pretty,” and sometimes “stupid.” For male characters, the words describing them included “leadership,” “power,” “strength,” and “politics.” The work shows that a data-driven application of multi-disciplinary methods from AI and social science can be used in support of literary culture and contribute to social understanding.

Contact: Mayank Kejriwal, [kejriwal@isi.edu](mailto:kejriwal@isi.edu)



The difference between male and female pronouns declines by more than half in female-authored books compared to male-authored books (top pane). Examples of adjectives associated with male characters versus female characters (bottom pane).

## Detecting Semantic Errors in Tables Using Textual Evidence

Web tables serve as a rich source of knowledge that supports many knowledge-driven intelligent applications. However, like other online resources, information in web tables is prone to errors and noise. Syntactic errors are erroneous values that relate to the representation or format of data and usually yield fatal errors when input into systems, databases, or applications. Common syntactic errors include missing values, typos, format inconsistencies, and violated attribute dependency. In contrast, semantic errors are attributed to the semantic meanings of the cell values and generally cannot be detected by traditional database systems.

The major challenge in detecting semantic errors is that external knowledge sources are required in order to validate the tables' values. Structured knowledge sources such as Wikidata have been used to extract essential information for detecting semantic errors. However, structured data sources require extensive manual effort to build and maintain, and thus they contain limited information compared to unstructured data.

To address the limited scope of data available from structured data sources, researchers in the Center on Knowledge Graphs exploit open-domain content such as Wikipedia as a source of evidential, textual data to verify the semantic correctness of tables.

The Semantic Error Detection (SEED) system has four phases:

- Document retrieval: selecting a set of relevant documents from Wikipedia using keyword search and dense representation scoring
- Sentence selection: segmenting all relevant documents into sentences
- Table verification: leveraging table linearization and sequence-to-sequence models to train a classification model to select evidential sentences
- Error cell correction: leveraging table linearization and sequence-to-sequence models to train a classification model to select evidential sentences

The same neural architecture is used in the table verification step to predict the correctness of table cell values. Erroneous tables are then input to a generic question-answering system for correction. Questions are generated based on tables' column header and values; contradicting sentences serve as contexts. The generated answers are the suggested corrections from the SEED system.

Contact: Craig Knoblock, [knoblock@isi.edu](mailto:knoblock@isi.edu)

*An erroneous table is input; the SEED system corrects it with extracted information from Wikipedia.*

Input Table			
<b>Title: Sam Underwood</b>			
Years	Title	Role	Location
2010	Equus	Alan Strang	HERE Arts Center
SEED System Processing			
Phase	Output		
Document retrieval	"Sam Underwood" "James Rado"		
Sentence selection	"[...], Underwood was asked to play the part of Alan Strang in a production of 'Equus' at the John Drew Theatre [...]"		
Table verification	"[...], Underwood was asked to play the part of Alan Strang in a production of 'Equus' at the John Drew Theatre [...]" ⇒ <b>Not entailed</b>		
Error cell correction	<b>Question:</b> Which location?  <b>Context:</b> "[...], Underwood was asked to play the part of Alan Strang in a production of 'Equus' at the John Drew Theatre [...]"  <b>Answer:</b> John Drew Theatre		

## Guidance to DNS Operators Builds on Years of Research

ISI operates one of the 13 systems providing service for the very top of the Internet’s naming system, the Domain Name System (DNS). The DNS translates human-understandable names like `isi.edu` into computer-usable information like Internet addresses, so it is part of every web page, email message, and almost all Internet uses.

Researchers in the Networking and Cybersecurity Division pair the operation of this worldwide, critical infrastructure with research conducted by computer scientists and graduate students. The results of these studies are shared through peer-reviewed publications and published Internet standards that guide the Internet industry.


In 2022, ISI researchers collaborated with researchers from SIDN Labs, the organization that runs country-level authoritative DNS service for the Netherlands. They co-authored an Internet standards document titled “Considerations for Large Authoritative DNS Server Operators.” This document, published as Internet Engineering Task Force’s RFC 9199 (Request for Comments), summarizes technical recommendations gained from multiple peer-reviewed academic papers as seen through the lens of the team’s operational experience. It builds on four years of work with the standards community to document networking best practices for operating DNS in combination with global IP anycast. IP anycast allows servers that are physically distributed around the world to provide greater capacity, resistance to denial-of-service attacks, and lower latency to users. IP anycast enables the 13 different systems in the DNS root to provide service from more than 1600 physical locations today.

The results and techniques of this research enable the USC ISI Root Server’s IP anycast to be optimized to balance traffic among its six globally distributed sites. The figure below vividly displays how the load shifts from site to site as activity increases and wanes based on each server’s local time. For example, the load peaks during the day in Europe (Amsterdam, the top dark red band) but peaks at other times in Asia (Singapore, the bottom green band) and the Americans (the middle four bands).



Scan QR code for more information.

Contact: [Wes Hardaker, hardaker@isi.edu](mailto:Wes Hardaker, hardaker@isi.edu)



The fraction of traffic to each of the six B-Root anycast sites over seven days in December 2022

## What Physical Changes Occur in the Brain When a Memory Is Made?

A multidisciplinary team of USC researchers drawing from ISI and the neuroscientists in the Dornsife School answered the question of how synapses change to encode memory in the brain by inducing a memory in larval zebrafish and then mapping changes in their transparent heads with brain cells lit up.

After six years of research, the researchers discovered that learning causes synapses to proliferate in some areas and disappear in others rather than merely changing their strength, as commonly thought. These changes in synapses may help explain how memories are formed and why certain types of memories are stronger than others. The study was published in the Proceedings of the National Academy of Sciences.

The study was made possible with a new type of cell labeling, and the custom-made microscope invented at USC. ISI researchers developed a cutting-edge way to extract information from the raw data and track and archive the data collected to make the findings as accessible and reproducible as possible. Before this work, it was only possible to determine a synapse's location in a living brain by modifying its structure and function, making comparisons before and after memory formation unfeasible.

The researchers were able to determine for the first time the strength and location of synapses before and after learning in the brain of a living zebrafish. Zebrafish are large enough to have brains that function like human brains but small and transparent enough to offer a window into the living brain. By keeping the intact fish alive, the researchers were able to compare synapses in the same brain over time, a breakthrough in the neuroscience field.

To create memories to measure, the research team had to develop new methods to induce a larval zebrafish to learn. They did this by training 12-day-old fish to associate a light turning on with being heated on the head with an infrared laser, an action they sought to avoid by attempting to swim away. Fish that learned to associate the light with the impending laser would flick their tails, indicating that they had learned. Five hours of training later, the team was able to observe and capture significant changes in these zebrafish brains. This enabled the specialized microscope to scan the brain and image where the synapses were located.


Knowing how synapses change during learning could have broad implications. For example, better understanding of aberrant memory formation that can lead to post-traumatic stress disorder or addiction could lead to novel approaches for treatment.

The result was hundreds of images and experiments that had to be processed and analyzed. Researchers in ISI's Informatics Systems Research Division developed innovative new algorithms that made this possible while keeping track of the large and complex experiments that were performed over the duration of the investigation.

The ISI researchers made the results of the investigation transparent and reproducible, with every piece of data searchable and available to any scientist at a public website, Mapping the Dynamic Synaptome. The ISI team integrated the Deriva platform for scientific data management into daily research activities. This

enabled an uncommonly high degree of transparency and reproducibility, with each figure directly supported by the data that went into its creation and all the raw microscope data, analyzed data, human interpretation of the data, analysis programs, and essential research resources available in the same structured, queryable, and downloadable form that was used to conduct the study. The comprehensive system designed for data sharing and analysis was useful during the experiments because the teams could access the data at all times; it will guide those who want to pursue this work in the future.

**Contact: Carl Kesselman, [carl@isi.edu](mailto:carl@isi.edu)**



*The synapses of a zebrafish brain are highlighted by a microscope; their location is identified using methods developed at ISI.*

## Improving Network Operations through Accurate Network State Discovery

Network operators require accurate information on the state and structure of their networks. Unfortunately, this information is traditionally kept in inventory, wiring, and configuration databases, which are often outdated. Only a small number of individuals have the relevant knowledge and skill to operate the network; replacing them is incredibly expensive.

ISI developed a series of novel techniques to rapidly determine the state and structure of a network from a *limited set of sensors*. These sensors leverage existing flows to collect data about the network structure and have three broad capabilities: topology extrapolation, topology interpolation, and state discovery.

Topology extrapolation determines the structure of the network outward from a sensor when no other sensor can reach the same network region. Because of this, extrapolating sensors have limited overall capability and are ideally suited to be deployed near network enclaves.

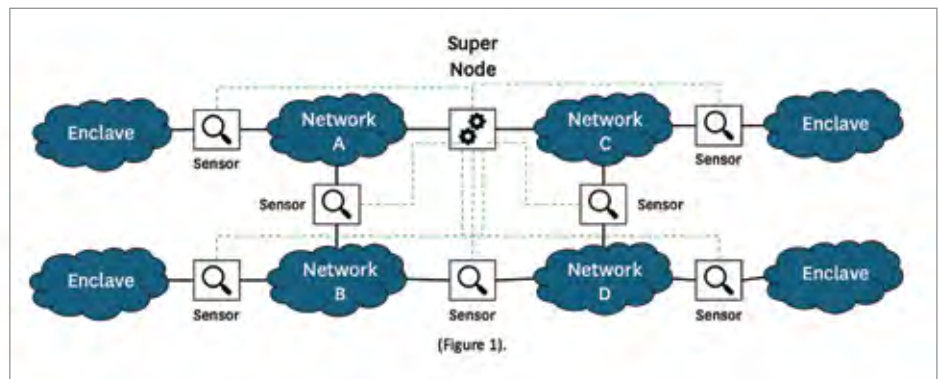
In contrast, topology interpolating sensors and state discovery capabilities leverage the fact that flows transit more than one sensor. Importantly, the sensors conduct most of their measurements by leveraging existing network traffic, making network overhead minimal. Lastly, state discovery is responsible for understanding the latencies, queuing status, and bottleneck bandwidths between network elements.

The information determined by the set of sensors is relayed to a super-node, whose job is to *infer* additional network state and structure from the combination of this information. A powerful capability is to infer the existence of hidden network structures (i.e., network elements that cannot be directly interacted with through sensor flows) based on how those hidden network structures interact with visible network traffic.

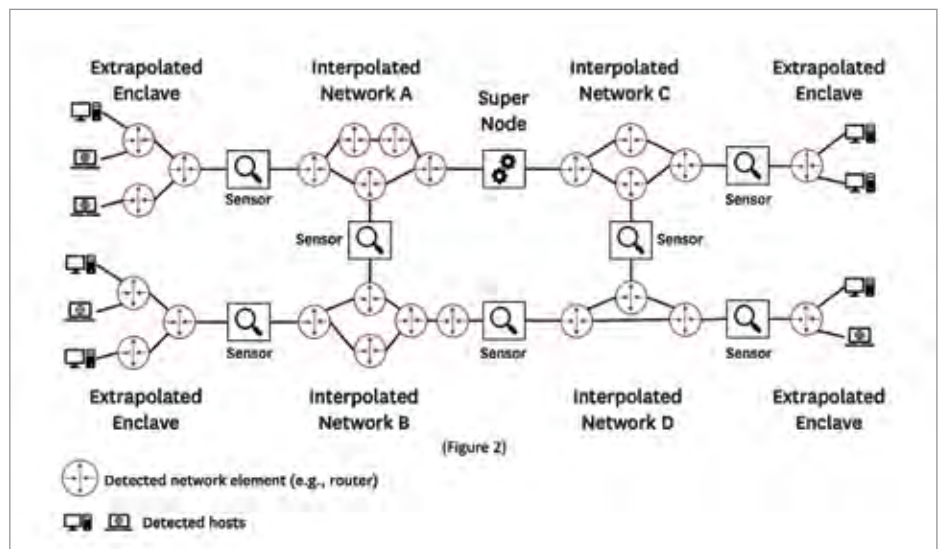
To reduce overall computational overhead, the super-nodes are responsible for calculating the underlying topology and storing this information. The super-nodes communicate with each other to form a consensus opinion regarding the entire topology based on their individual local information. Determining which nodes should be super-nodes is accomplished by solving the vertex k-center problem, an NP-hard problem but readily solvable through heuristic techniques.

This research generated strong results, with the demonstrated ability to scale to topologies of over 1000 nodes with only 48 sensors. The accuracy remains high for all tested topologies while maintaining very low overhead.

Contact: Erik Kline, [kline@isi.edu](mailto:kline@isi.edu)



Deployment of sensors and super nodes in a network. Interpolation occurs between sensors; extrapolation occurs from one sensor (towards an enclave). All information is sent to the super-node where inference occurs.



Inferred topology based on measurements from sensors



## SACHET: Finding Risky Code Submissions in Open Source Software

In large open source software projects open to contributions from any developer, such as the Linux kernel, code quality is maintained by trusted developers and reviewers. These projects are vulnerable to inadvertent errors or sociotechnical attacks. For example, a large number of inconsequential changes may be sent in a small region of the code, increasing the cognitive load for the reviewers of that region so that a vulnerability can be inserted without detection. It is not feasible to check all the contributed code for errors due to the size and complexity of the project. Still, social activity related to the code can focus attention on the areas of potential sociotechnical attacks.

Researchers in the Networking and Cybersecurity Division developed the SACHET dashboard, which combines a set of detectors for potentially risky code or social activity to assist an analyst in protecting open source code. These detectors are developed using machine learning to recognize situations associated with risky code submissions, combining three kinds of data:

- The contributed code
- Discussion in mailing lists, including an emotion analysis of comments and messages based on off-the-shelf software
- Typical behavior of developers and reviewers

For example, a classifier is trained to predict when a developer will propose their first patch-altering cryptographic code, based on their past behavior. A detector based on this classifier warns the analyst when a proposed change to cryptographic code is from a surprising source.

The SACHET dashboard brings information together from a range of detectors. It has surfaced a number of high-risk cases, including developers who are independently known for controversial interactions with the community and hidden security fixes. Hidden security fixes may be dangerous since they may alert attackers to vulnerabilities, but users are not urged to apply the patch.

**Contact: Jim Blythe, [blythe@isi.edu](mailto:blythe@isi.edu)**



Part of the SACHET dashboard, this view shows recent activity in the subtopic of Linux Secure Modules (LSM). The bars show the total number of messages in a day, and the dots below are color-coded to individual developers who submit patches. The lines above show changes from the previous five patches, of authors (green) and of files affected (purple). The bursts in activity that coincide with such a change in focus in June and July show the introduction of the new AppArmor module.

## Auditing Algorithmic Fairness in Job Ads

Ad platforms such as Facebook, Google, and LinkedIn promise value for advertisers through their targeted advertising. Unfortunately, these platforms often select ads skewed by gender or race. Social media platforms select ads using relevance algorithms that are proprietary to the platform.

A Networking and Cybersecurity Division researcher, working with a student and Princeton professor, demonstrated that external auditors can evaluate job ad skew, even while accounting for legitimate differences in applicant qualifications.

They carried out a study to evaluate Facebook and LinkedIn for potential bias in ad delivery. The study needed to isolate the platform’s algorithms from many other confounding factors (individual qualifications, number of samples, etc.). They compared Facebook and LinkedIn by running paired ads (ads for similar jobs at companies with different reputations) against populations of known gender. When the gender response between the pair of ads differed by a statistically significant amount, they concluded there was evidence that the platform was skewing ad impressions toward that gender.

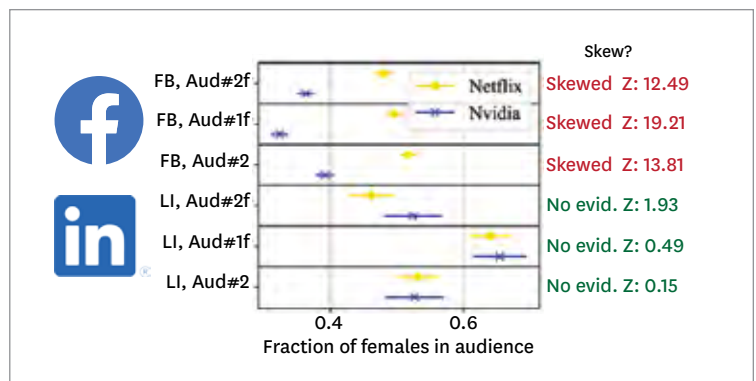
The figure shows experiments over three populations for Facebook and LinkedIn. It shows that Facebook displays some ads to fewer women since the blue and yellow dots are statistically different; their error bars do not overlap, and their Z values are large. LinkedIn does not show evidence of algorithmic skew; their ads are shown to similar fractions of each gender since the blue and yellow dots are within error bars and the Z values are below a threshold.

The results demonstrated a statistically significant skew for Facebook’s job ads and showed an absence of skew on LinkedIn.

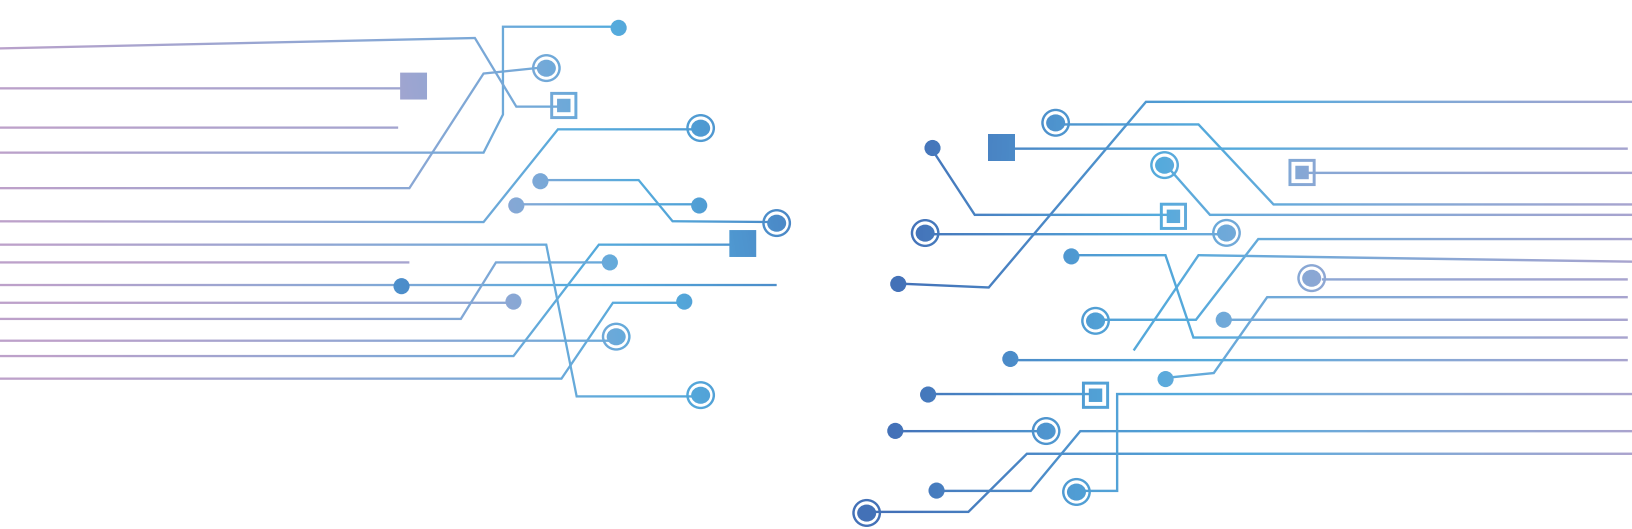


Scan QR code for more information.

Contact: John Heidemann, johnh@isi.edu



Evidence of skew comparing gender distributions in ads for software engineering comparing Facebook (FB) and LinkedIn (LI) using paired ads (for software engineering jobs at Netflix, top yellow, and Nvidia, bottom blue). Different results for each pair—the yellow and blue dots differ, and the Z values are large—suggest that Facebook’s algorithms show more Nvidia ads to men than expected, while LinkedIn shows no evidence of skew.



## Harvesting Planck Radiation for Optical Communication

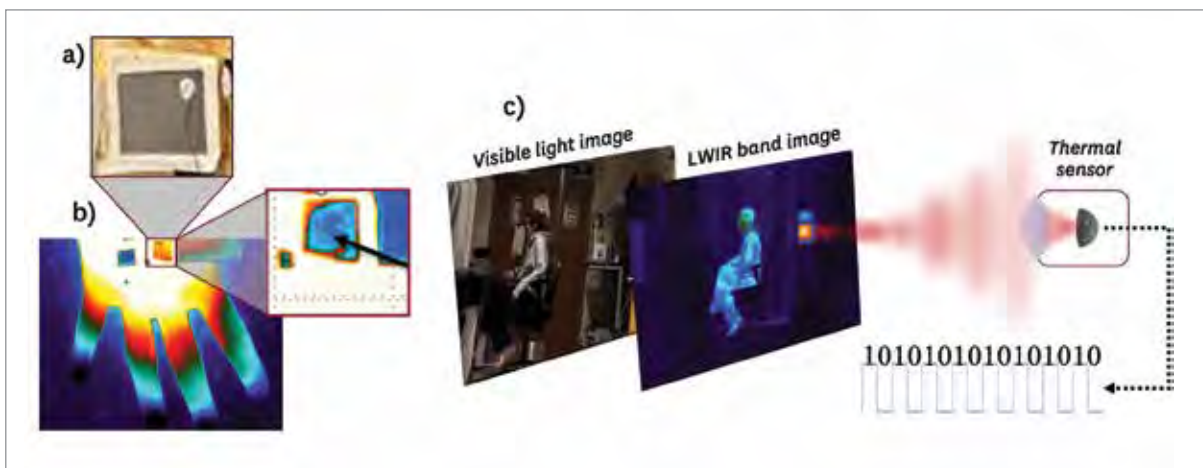
All warm objects generate and emit a well-understood type of electromagnetic emissions, or heat energy, called Planck radiation. Although invisible to the naked eye, Planck radiation is an intense source of energy that radiates persistently. As an example, the average human being generates approximately 500 watts of Planck radiation at all times—and more when exercising. This Planck radiation is considered “waste heat” and is expelled from our bodies to keep ourselves cool.

In the Laboratory for Quantum-Limited Information, researchers developed a new type of material named a “Planck modulator” that can harvest this Planck radiation and use it to implement a revolutionary communications paradigm.

The Planck modulator shown in panel a) is a compound (multi-layer) device made from graphene. By applying an electrical voltage to the Planck modulator, the researchers can change the intensity of the emitted Planck radiation, shown by the apparent “cooling” of the device from a “hot” (red) to “cold” (blue) shown in panel b).

Illustrated in panel c), the team used this capability to modulate Planck radiation intensity, demonstrating a proof-of-concept communications channel operating at 100 bits per second using Planck radiation similar to that generated by a human—the first demonstration of its kind. Harvesting heat to enable communications could enable new communications modalities for search and rescue missions using heat from a person to power a communications transmitter to signal information about location or health. Additionally, the team is investigating using this technology for provably private communications, hiding communications in the omnipresent thermal noise.

**Contact: Jonathan Lenahan Habif, [habif@isi.edu](mailto:habif@isi.edu)**



a) Planck modulator is a multi-layer device fabricated from graphene, an exotic formation of carbon atoms. b) Planck modulator is shown resting on an experimenter’s hand. When an electrical voltage is applied to the Planck modulator, it reduces the intensity of the Planck radiation emitted, shown as an apparent “cooling” of the device from red to blue (inset). c) Using the functionality of the Planck modulator, the team performed a first-ever optical communications demonstration harvesting heat energy and using it to encode information.

## Exploring the Information Ecosystem with News Article Revision Histories

News articles are often updated after initial publication, and these revisions can provide rich insights into the journalistic process. Which facts in a news article are most likely to change? How are quotes fact-checked and contextualized? What kinds of information are added after the initial story breaks? Stylistic changes? Analysis? Background?

News article revision histories provide clues to narrative and factual evolution in news articles. Articles may be revised to incorporate new information, update events, and make stylistic changes.

Researchers in the Artificial Intelligence Division built the first publicly available dataset of news revision histories, NewsEdits, annotated with sentence-level semantically coherent change information. The massive dataset of news revision histories contains 1.2 million articles with 4.6 million versions from over 22 English- and French-language newspaper sources based in three countries. The articles span 15 years of coverage (2006-2021).

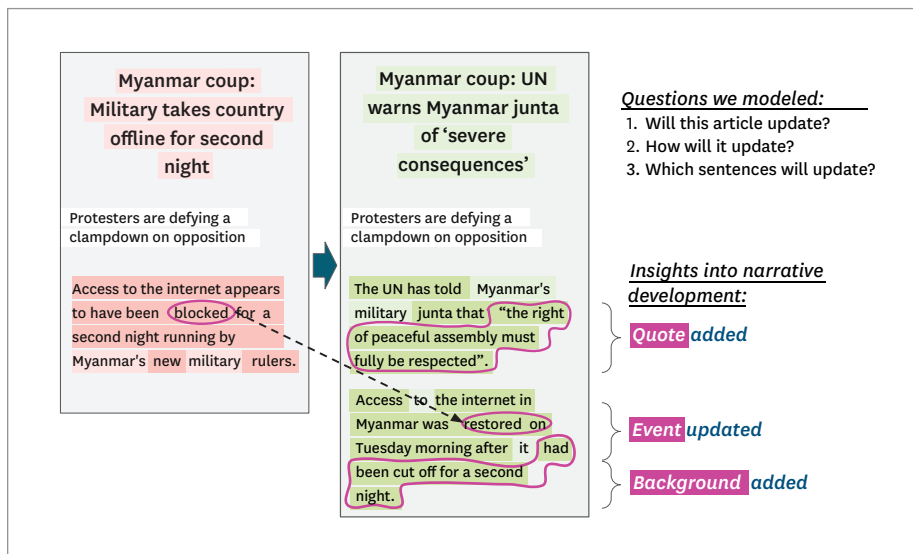
The team developed a document-level edits model that represents revisions as sequences of syntactic operations in terms of Sentence Deletions, Sentence Additions, and Sentence Updates, which they used to analyze revisions. They developed a high-accuracy extraction algorithm to process articles and identify these sentence-level actions. The researchers learned that many edits are made to update events, incorporate additional quotes, and provide main content.

To underscore the factual nature of many edit actions, they conducted analyses showing that added and deleted sentences are more likely to contain revised events, main content, and quotes than updated sentences. To explore whether these actions are predictable, the team introduced three novel tasks to predict the type and location of an action that would lead to a new revision. They learned that experts could, in fact, predict edit actions but that this is challenging for large language models.



Scan QR code for more information.

Contact: Alexander Spangher, [spangher@isi.edu](mailto:spangher@isi.edu)



Scalable algorithms identify if sentences were deleted, added, or updated between news article revisions. Edits to news articles are found to be factual in nature: quotes, main events, and event updates. Furthermore, models can predict the way breaking news will update.

## Workflows Support USC's New Cryo-Electron Microscopes

ISI researchers worked with USC's Center for Advanced Research Computing (CARC) to develop a solution for preprocessing cryo-electron microscopy (cryo-EM) data, making it more efficient for scientists to manage the huge datasets generated by USC's cryo-electron microscopes.

USC's cryogenic electron microscopy facility officially opened in 2022. The cryogenic electron microscopy facility houses two state-of-the-art electron microscopes capable of imaging molecules. Research with these instruments focuses on structural biology and studying the shapes of biological components, such as proteins and ribosomes, and how their shapes change as they perform their tasks within the cell. This area is important for understanding how molecules function within cells and for developing new therapies.

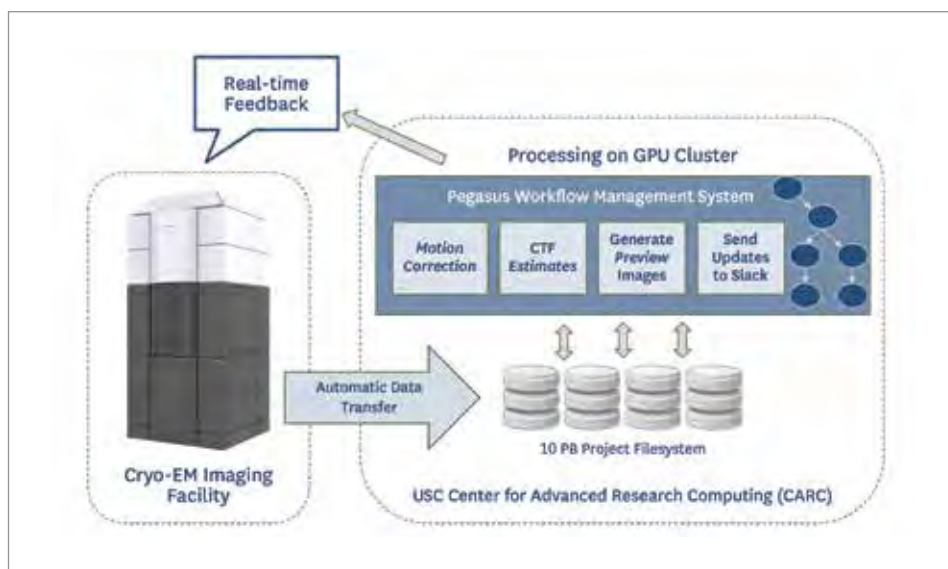
Unlike traditional microscopes, where a researcher looks into the telescope and sees magnified samples, electron microscopes generate terabytes of data that must be processed and visualized in order to be seen and inspected. The facility housing the electron microscopes does not have the capacity to process the data generated by the electron microscopes; the control PCs are in a separate, dedicated facility.

To make these microscopes widely available, the research community at the USC Core Center of Excellence in Nano Imaging approached USC's Center for Advanced Research Computing (CARC) to coordinate the development of an automated solution for transferring and processing the experimental data. CARC operates a high performance computing (HPC) cluster suited for processing these large datasets.

A key aspect of electron microscopy is the fine-tuning of data acquisition, where the preprocessing of data is performed in near real-time at lower magnification, allowing the researcher to decide which grids of the sample need full data acquisition. This is important in the context of the storage and computing resources required for processing the full dataset.

CARC leadership contacted researchers in ISI's Science Automation Technologies group for assistance in developing an automated solution. Working with CARC staff, the ISI researchers developed a lightweight automated Cryo-EM image pre-processing service that keeps user interaction minimal. Users have the option to start the pre-processing immediately after initiating the microscope session. They receive real-time feedback that enables them to adjust parameters during data acquisition. The service manages the data transfers required for preprocessing and efficiently launches the tasks on the CARC HPC resources to optimize the storage footprint and overall execution time of the workflows.

**Contact:** Ewa Deelman, [deelman@isi.edu](mailto:deelman@isi.edu)



*Workflow for USC's cryogenic electron microscopes*

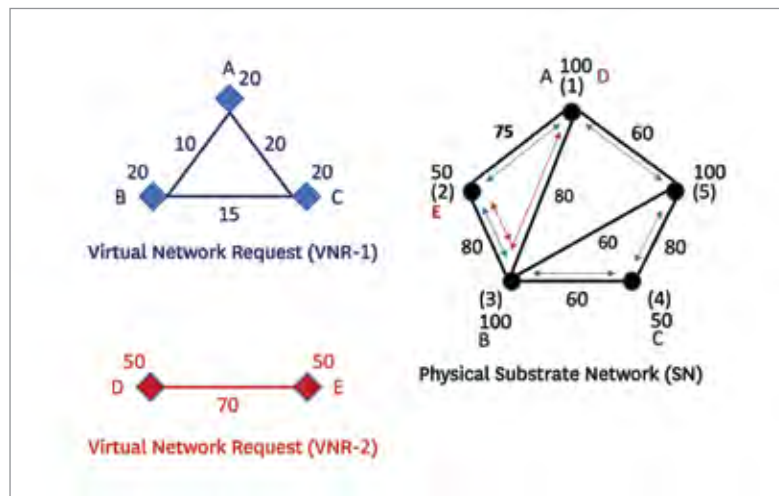
## Applying Combinatorial Optimization to Network Embedding Problems

Traditionally, mobile networks have employed a “one-size-fits-all” paradigm for deploying network resources to support different applications. This model is not sufficient for addressing the challenges of a heterogeneous market that supports a vast number of different kinds of applications, such as secure real-time communication, video content delivery, and geolocation. Network slicing in 5G networks refers to technologies that can accommodate such a heterogeneous market by properly managing the resources of a physical substrate network (SN). The idea is to allocate communication paths in the SN to implement logical connections needed for various service requirements. Network orchestration is achieved by coordinating the resources allocated for each network slice to fully exploit the SN’s resources and maximize the revenue of the network provider.

At the core of network slicing is the combinatorial problem of virtual network embedding (VNE). Each SN vertex has a CPU resource, and each SN edge has a bandwidth resource. Virtual network requests (VNRs), which are essentially requests for resources, must be satisfied. Each VNR vertex has a CPU requirement, and each VNR edge has a bandwidth requirement. Each VNR vertex must be mapped to an SN vertex. Each VNR edge must be mapped to an SN path such that: a) every SN vertex can meet the CPU requirements of all VNR vertices assigned to it; b) every SN edge can meet the bandwidth requirements of all VNR edges that utilize it; and c) the total cost of SN resources utilized for the mapping is minimized. Additional mapping constraints might also exist. The figure illustrates the VNE problem.

Computing an optimal solution to the VNE problem is NP-hard. ISI researchers have applied their experience in combinatorial optimization to this problem domain. Their novel algorithmic framework, conflict-based search (CBS), is currently the best approach for solving the VNE problem, outperforming competing state-of-the-art methods in both efficiency and quality of the solutions produced. In general, the conflict-based search approach is applicable whenever resources on a network must be managed efficiently and cost-effectively.

Contact: Satish Kumar Thittamarahalli, [skumar@isi.edu](mailto:skumar@isi.edu)



Two VNRs, VNR-1 (blue) and VNR-2 (red), are on the left. The SN (black) is on the right. The non-negative numbers annotating the VNR/SN vertices (edges) represent their CPU (bandwidth) requirements/capacities. A VNE solution, mapping each VNR vertex to an SN vertex and each VNR edge to an SN path, is shown by the solid arrows.

ISI researchers were instrumental in working with the USC Computer Science Department leadership to create USC’s Data Science Program in the Viterbi School of Engineering. Keston Executive Director Craig Knoblock was the first director of the program.

Today the program is thriving, and with more than a thousand undergraduate and graduate students, it is one of the largest programs in Viterbi. The Data Science Program offers these degrees:

- Bachelor of Arts in Data Science
- Undergraduate minor in Foundations of Data Science
- Bachelor of Science in Artificial Intelligence for Business (joint with the USC Marshall School of Business)
- Master of Science in Applied Data Science
- Master of Science in Spatial Data Science (joint with the USC Dornsife College of Letters, Arts, and Sciences)
- Master of Science in Communication Data Science
- Master of Science in Cyber Security Engineering
- Master of Science in Healthcare Data Science
- Master of Science in Communication Data Science (dual degree with Tsinghua University)
- Master of Science in Environmental Data Science (joint with the USC Dornsife College of Letters, Arts, and Sciences)
- Master of Science in Public Policy Data Science (joint with the USC Price School of Public Policy)
- Graduate Certificate in Data Science Foundations
- Graduate Certificate in Applied Data Science
- Progressive Degree Program

The interdisciplinary data science degrees are the first of their kind at USC and are designed to introduce non-computer science majors to careers in data science.

**DATA SCIENCE PROGRAM LEADERSHIP AND FACULTY**

ISI’s **Yolanda Gil** is Director of the Data Science Program; **Fred Morstatter** is Associate Director. ISI researchers who work actively in data science research teach in the Viterbi Data Science Program. They have designed introductory and advanced courses in scalable data systems, machine learning, knowledge graphs, and AI and data ethics.

- Marcin Abram, Lecturer*
- Jeremy Abramson, Lecturer*
- Jose Luis Ambite, Research Associate Professor*
- Yigal Arens, Research Professor*
- Jim Blythe, Lecturer*
- Keith Burghardt, Lecturer*
- Ulf Hermjakob, Lecturer*
- Filip Ilievski, Research Assistant Professor*
- Carl Kesselman, Professor*

- Deborah Khider, Lecturer*
- Kristina Lerman, Research Professor*
- Fred Morstatter, Research Assistant Professor*
- Jay Pujara, Research Assistant Professor*
- Mohammad Rostami, Research Assistant Professor*
- Gleb Satyukov, Lecturer*
- Satish Thittamaranahalli, Research Assistant Professor*
- Ke-Thia Yao, Lecturer*



*Datafest 2022*

## NEW SPONSORED RESEARCH AWARDS

In 2022, ISI received 54 new funded research awards. Sponsored by the federal government, industry, and philanthropy, the awards range from tens of thousands of dollars to millions. The basic and applied research supported by these awards spans many areas: quantum information science, machine learning, microelectronics security, quantifying and analyzing synapse structure in zebrafish brains, and more.

### AIR FORCE

#### **Modeling and Simulation for Flexible structures for Space Debris Removal**

**David Barnhart**, *Principal Investigator*

(via L.Garde. Inc.)

#### **Transformational 3D+ Scene Reconstruction Platform for Advanced OSAM Proximity Operations**

**David Barnhart**, *Principal Investigator*

(via Quidient, LLC.)

### AIR FORCE OFFICE OF SCIENTIFIC RESEARCH

#### **Assured Replaceable Components Enabling Extension-of-Life (ARCEE)**

**Alefiya Hussain**, *Principal Investigator*

(via Tangram Flex, Inc.)

#### **Modernizing DeterLab Education Support and Growing and Diversifying Education User Base**

**Jelena Mirkovic**, *Principal Investigator*

#### **Learning Dynamics and Detecting Causal Pathways In Coupled Online-Offline Systems**

**Fred Morstatter**, *Principal Investigator*

(via UCLA)

#### **SOUL ADR Demonstration Support**

**David Barnhart**, *Principal Investigator*

(via Busek Co., Inc.)

### AIR FORCE RESEARCH LABORATORY

#### **Information Extraction for New Emerging Noisy User-Generated Micro-Text**

**Elizabeth Boschee**, *Principal Investigator*

(via InferLink Corporation)

#### **DISCOVER: A Data-driven Integrated Approach for Semantic Inconsistencies Verification**

**Wael AbdAlmageed**, *Principal Investigator*

(via Purdue University)

#### **Unprepared Surface Capture with Multiple Arm Architecture Analysis: Extending REACCH**

**David Barnhart**, *Principal Investigator*

(via Kall Morris, Inc.)

### ALFRED P. SLOAN FOUNDATION

#### **Privacy Preserving Entity Resolution**

**Srivatsan Ravi**, *Principal Investigator*

(via Actuate Innovation)

### BILL & MELINDA GATES FOUNDATION

#### **Handwriting Recognition in the Decennial Census Digitization and Linkage Project**

**Huaigu Cao**, *Principal Investigator*

(via University of Michigan)



## DEFENSE ADVANCED PROJECTS AGENCY (DARPA)

### DANUBE

**Stephen Schwab**, *Principal Investigator*  
(via University of Maryland)

### ALLAN: Agents Learning Lying and Negotiation

**Jonathan May**, *Principal Investigator*  
(via University of Maryland)

### EMPATH: Predicting Emergent Pathways of Information and Influence

**Kristina Lerman**, *Principal Investigator*

### Tailored INCAS EcoSystems TIES

**Wael AbdAlmageed**, *Principal Investigator*  
(via Lockheed Martin)

### MIDAG: Modeling Influence Pathways with Multi-Dimensional Dynamic Graphs

**Xuezhe Ma**, *Principal Investigator*  
(via UCLA)

### BeQuEST: Benchmarking Quantum Enhancement in Science & Technology

**Itay Hen**, *Principal Investigator*

### KNIC: Knowledge Needed In Context

**Jay Pujara**, *Principal Investigator*

### DARPA AIE Civil Sanctuary

**Fred Morstatter**, *Principal Investigator*  
(via Aptima, Inc.)

### DIAMOND: All-Optical Computing Directly on Optical Fiber Data Traffic

**Jonathan Habif**, *Principal Investigator*

## DEPARTMENT OF DEFENSE

### MA - Intergovernmental Personnel Act (IPA)

**Terry Benzel**, *Principal Investigator*

### Galactic Gopher

**Travis Haroldsen**, *Principal Investigator*  
(via Georgia Institute of Technology)

## DEPARTMENT OF ENERGY

### BssW Fellowship

**Ewa Deelman**, *Principal Investigator*  
(via Krell Institute)

## DEPARTMENT OF HOMELAND SECURITY

### CRATES Support, ITSME Support

**David Balenson**, *Principal Investigator*  
(via SRI International)

## INTELLIGENCE ADVANCED RESEARCH PROJECTS ACTIVITY

### Stylometric Authorship Discernment and Interpretation for Realistic Inputs

**Elizabeth Boschee**, *Principal Investigator*

## NEW SPONSORED RESEARCH AWARDS (CONTINUED)

### JET PROPULSION LABORATORY

Software Fault Tolerance, Design and Development

John Paul Walters, *Principal Investigator*

(via Microchip Technology, Inc.)

### JOHN TEMPLETON FOUNDATION

Maryland Quantum-Thermodynamics Hub: Decoherence and Equilibration in Quantum Systems

Amir Kalev, *Principal Investigator*

(via University of Maryland)

### KAISER PERMANENTE

Healthcare Knowledge Graph for Search and Recommender Systems

Mayank Kejriwal, *Principal Investigator*

### NATIONAL SCIENCE FOUNDATION

Conference: Incorporating Ethics into the Human-Centered Design of AI Solutions

Michael Pazzani, *Principal Investigator*

(via University of Maryland)

Collaborative Research: SII-NRDZ: ASPIRE: Advanced Spectrum Initiative for Research and Experimentation

Alefiya Hussain, *Principal Investigator*

Collaborative Research: CCRI: NEW: Open Community Platform for Sharing Vehicle Telematics Data for Research and Innovation

David Balenson, *Principal Investigator*

SII-NRDZ: STRATA - Study of INL for Radio National Facility

Erik Kline, *Principal Investigator*

Collaborative Research: CNS Core: Medium: A Traffic Map for the Internet

John Heidemann, *Principal Investigator*

IMR: RI-P: Safe and Flexible Experimental Dataset Access and Sharing-Planning (SAFED-ASP)

John Heidemann, *Principal Investigator*

Collaborative Research: FMitF: Track I: A Formal Verification and Implementation Stack for Programmable Logic Controllers

Luis Garcia, *Principal Investigator*

CRII: III: Robust and Explainable AI Agents with Common Sense

Filip Ilievski, *Principal Investigator*

PIPP Phase I: Evaluating the Effectiveness of Messaging and Modeling During Pandemics (PandEval)

Kristina Lerman, *Principal Investigator*

(via University of Maryland)

NSF-BSF: Fast Quantum Optimal Control on Exponentially Large Spaces

Amir Kalev, *Principal Investigator*

Track 2: Customized Multi-tier Assistance, Training, and Computational Help (MATCH) for End User ACCESS to CI

Ewa Deelman, *Principal Investigator*

(via University of Colorado-Boulder)

CC\* Regional Computing: Building Cyberinfrastructure to Forge a Regional Research Computing Alliance in Southern California

Carl Kesselman, *Principal Investigator*

## NATIONAL GEOSPATIAL-INTELLIGENCE AGENCY

Combining Efficient Algorithms, Machine Learning and Knowledge Graphs for Scalable, High-dimensional Nearest Neighbor Search

Wael AbdAlmageed, *Principal Investigator*

## NATIONAL INSTITUTES OF HEALTH

CRCNS Data Sharing Proposal: Collaborative Research: NeuroBridge: Connecting Big Data for Reproducibility Research in Clinical Neuroscience

Jose Luis Ambite, *Principal Investigator*

(via Ohio State University)

SA - ISI-Horn-Analyzing Digital Menu Data to Characterize Nutritional Quality of Food Environments in Latino Neighborhoods within the City of Los Angeles

Keith Burghardt, *Principal Investigator*

(via Children's Hospital Los Angeles)

## NATIONAL RECONNAISSANCE OFFICE

Quantum-Optimal Frequency Estimation Program

Jonathan Habif, *Principal Investigator*

(via BAE Systems)

## OFFICE OF NAVAL RESEARCH

Graphene-Based Materials for Long-Wave Infrared Emission Modulation

Jonathan Habif, *Principal Investigator*

## OPEN PHILANTHROPY

iPASSAGE: improved Performance, Analytics and Summarization of Synergistic Anticipation of Geopolitical Events

Fred Morstatter, *Principal Investigator*

(via University of Maryland)

Developing Artificial Intelligence Forecasting Questions

Mayank Kejriwal, *Principal Investigator*

(via University of Maryland)

## SERVEL GLOBAL SOLUTIONS, LLC

Measure and Recommend Tools for Increasing Speed and Improving Quality of Bible Translation

Ulf Hermjakob, *Principal Investigator*

## SWISS FEDERAL DEPARTMENT OF DEFENCE

Knowledge Graph Technology for Information Analytics

Filip Ilievski, *Principal Investigator*

## US GOVERNMENT

ADES - Advanced Defense Evaluation Through Simulation

Michael Collins, *Principal Investigator*

Continual Learning of Few-Shot Learners for Natural Language Processing

Mohammad Rostami, *Principal Investigator*

(via North Carolina State University)

SEARCHER Improvements for JHU COE

Scott Miller, *Principal Investigator*

(via Johns Hopkins University)

## WYCLIFFE BIBLE TRANSLATORS

Greek Room: Bible Translation Checking and First Draft Translation Support

Ulf Hermjakob, *Principal Investigator*

## ISI Fellows

The Fellow is the highest scientific/engineering rank of ISI and is held by only a small fraction of the research staff. The researchers who achieve this rank have a widely recognized international reputation. ISI Fellows have demonstrated exceptional achievements in science or engineering relevant to the institute's core missions, or they have provided critical technical leadership of major projects impacting ISI's future directions or its mission. In 2022, these researchers were named ISI Fellows.



### Ewa Deelman

Research Professor  
Research Director  
Principal Scientist

For contributions to the automation of scientific computing applications and advancing the role of computation in scientific domains of societal importance.



### Yolanda Gil

Research Professor  
Viterbi Director for Major Strategic AI and Data Science Initiatives  
Director of Knowledge Technologies  
Principal Scientist

For contributions on knowledge capture and scientific workflows and for promoting ISI's artificial intelligence research to the larger community.

## Professional and Academic Recognition

### Wael AbdAlmageed

#### Invited Talk

*End-to-end Biometrics: Yesterday, Today, and Tomorrow*  
Biometrics Institute

#### Invited Talk

*Fake Visual Evidence—Creation, Detection, and Admissibility*  
European Commission

#### Invited Talk

*If Not Malicious, Then Just for The Thrill of It*  
Amazon One Virtual

#### Invited Talk

*If Not Malicious, Then Just for the Thrill of It*  
University of Virginia Law

#### Invited Talk

*Combatting Misinformation*  
Renaissance Weekend

#### Invited Talk

*The Dangers of Deepfakes and AI-generated Media*  
USC Annenberg School for Communication and Journalism

#### Keynote

*Biometrics Under Attack—Where Do We Go from Here?*  
Workshop On Manipulation, Adversarial, and Presentation Attacks in Biometrics (Map-A)  
Winter Conference on Computer Vision

**ISI Achievement Award** - for founding and directing the Visual Intelligence and Multimedia Analytics Lab and his research on deep fake detection

**Peter A. Beerel**

**Best Paper Award Nomination**  
IFIP/IEEE VLSI-SoC Conference 2022

**Terry Benzel**

**Outstanding Service Award**  
IEEE Security & Privacy Magazine

**Recognition of Service**  
Office of the Director, National Science Foundation

**Golden Core Member**  
IEEE Computer Society

**Keith Burghardt**

**Runner-Up Best Paper Award (Research Track)**  
2022 IEEE/ACM International Conference on Advances in Social Network Analysis and Mining

**Muhao Chen**

**Amazon Research Award**

**Tainā Coleman**

**Best Research Assistant Award**  
USC Viterbi Department of Computer Science

**Ewa Deelman**

**Euro-Par Achievement Award**  
Euro-Par Conference

**Keynote**  
*The Pegasus Workflow Management System: Using Clouds and More for Science*  
10th IEEE International Conference on Cloud Engineering

**Keynote**  
*Workflows for Science: How Workflows Need to Adapt to Modern Application Needs*  
CHPC National Conference 2022

**Weiwei Duan, Fandel Lin, Zekun Li**

**First Place**  
Map Feature Extraction Challenge  
DARPA, US Geological Survey

**Katy Felkner**

**NSF Graduate Research Fellowship (GRFP)**

**Luis Garcia, Christophe Hauser, and Haoda Wang**

**Best Paper Award**  
*AutoCPS: Control Software Dataset Generation for Semantic Reverse Engineering*  
IEEE Workshop on the Internet of Safe Things

**Tanmay Ghai**

**2022 Viterbi Master's Student Award for Best Research in the Computer Science Department**

**Yolanda Gil**

**M. Lee Allison Award for Outstanding Contributions to Geoinformatics**  
Geological Society of America

**John Heidemann et al.**

**Best Paper Award**  
*Old but Gold: Prospecting TCP to Engineer and Live Monitor DNS Anycast*  
Passive and Active Measurement Conference (PAM 2022)

**Abigail Horn**

**Invited Symposium Lecture**  
*Opportunities for Big Data and Mathematical Models of the Food Supply System for Food Safety Risk Assessment*  
2022 Nanoscale Science and Engineering for Agriculture and Food Systems Gordon Research Conference

**Abigail Horn, Keith Burghardt, Alex Seo, and Andres Abeliuk**

**Top Poster Award**  
Health Equity in Action Annual Workshop of the NIH National Institute on Minority Health and Health Disparities (NIMHD) Research Coordinating Center to Reduce Disparities in Multiple Chronic Diseases (RCC-RD-MCD)

**Filip Ilievski**

**Keynote**  
*Commonsense Knowledge and Reasoning Workshop on Reasoning with Imperfect Knowledge*  
Knowledge Graph Conference (KGC)

**Invited Participant and Session Chair**  
*Knowledge Graphs and their Role in the Knowledge Engineering of the 21st Century*  
Dagstuhl Seminar: Tools and User Experience for Knowledge Graph Engineering

**Akhilesh Jaiswal**

**IEEE Brain Community Best Paper Award**  
*Neural Computing with Magnetolectric Domain-Wall-Based Neurosynaptic Devices*

**Julie Jiang**

**Snap Research Fellowship**  
Snap Inc.

## HONORS AND AWARDS (CONTINUED)

### Emmanuel Johnson

**New Brunswick High School Hall of Distinguish Alumni**  
New Brunswick Education Foundation

**The Order of Arête**  
University of Southern California

**TIER-ED Speaker Series Invited Speaker**  
University of Illinois Urbana-Champaign

**Invited Participant**  
Heidelberg Laureate Forum

### Mayank Kejriwal

**Mike Sargeant Medal for Early Career Achievement**  
Institution of Engineering and Technology (IET)

**Science Policy Fellowship**  
Society for Industrial and Applied Mathematics (SIAM)

**Keynote**  
*Domain-specific Knowledge Graph Construction: Opportunities and Challenges*  
IEEE ICDM 2022 International Workshop on Knowledge Graphs

**40 Under 40 in 2022**  
CDO Magazine

**Carl Kesselman and Ian Foster (University of Chicago)**  
IEEE Internet Award (2023)

**Carl Kesselman**  
IEEE Fellow (2023)

**Deborah Khider**  
**Community Service and Leadership Award**  
EarthCube

**Member**  
Science, User, and Facility Advisory Board  
Earth System Grid Federation

### Kristina Lerman

**Keynote**  
*Biases in Data & Other Threats to Validity of Predictive Models*  
WebMedia 2022

**AAAI Fellow**  
American Association for the Advancement of Artificial Intelligence

**Computer Science Distinguished Lecture**  
*Bias and Emergent Instabilities in Socially-Embedded Algorithms*  
University of Illinois Chicago

**Keynote**  
*The Emergence of Gender Disparities and the Glass Ceiling Effect in Science*  
NetSci 2022 Satellite on Network Inequalities

### Jonathan May

**Invited Talk**  
*Theory of Mind and Intent in Procedural Agents: Learning Lessons from Improvisational Theater and Text-based Adventure Games*  
Computer and Information Science Department,  
University of Pennsylvania

### Jonathan May, Alexander Spangher, Xiang Ren, and Nanyun Peng (UCLA)

**Outstanding Paper Award**  
*NewsEdits: A Dataset of News Article Revision Histories and a Novel Document-Level Reasoning Challenge*  
North American Chapter of the Association for Computational Linguistics (NAACL)

### Shrikanth Narayanan

**Guggenheim Fellowship**  
John Simon Guggenheim Foundation

**Fellow**  
Association for the Advancement of Affective Computing (AAAC)

**Plenary Lecture**  
30th Brazilian Congress on Speech and Language Pathology

**Keynote**  
International Conference on Imaging, Signal Processing and Communications (ICISPC)

**Invited Talk**  
*Emotions in Organizational and Social Media Communication*  
International Society for Research on Emotion (ISRE) Workshop

**Keynote**  
*CLPsych: Workshop on Computational Linguistics and Clinical Psychology*  
North American Chapter of the Association for Computational Linguistics (NAACL)

**Invited Talk**

Society for Psychotherapy Research (SPR) International Annual Meeting

**Invited Talk**

Expert Session, IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)

**Plenary Lecture**

AAAI Spring Symposium: Designing AI for Open Worlds  
American Association for the Advancement of Artificial Intelligence

**Invited Talk**

University of Miami Institute for Data Science & Computing

**ASM Rizvi et al.**

**Best Paper Award**

*Defending Root DNS Servers Against DDoS Using Layered Defenses*

IEEE International Conference on Communications Systems and Networks (COMSNETS)

**Mohammad Rostami**

**AAAI New Faculty Highlights**

**Ganesh Chennimalai Sankaran**

**Invited Talk**

*Cooperative Network Functions*

ACM SIGMETRICS Workshop on Performance of Host-based Network Applications (PerfNA 2022)

**Basel Shbita**

**2022 University Outstanding Teaching Assistant Award**

University of Southern California

**Leonidas Spinoulas**

**Invited Presentation**

*Multispectral Biometrics System Framework: Application to Presentation Attack Detection*

IEEE Sensors Conference

**Karan Vahi**

**2022 Better Scientific Software (BSSw) Fellow**

Department of Energy and National Sciences Foundation

**Haoda Wang**

**2022 Best Poster Award of ECE 12th Annual Research Festival**

USC Viterbi Ming Hsieh Department of Electrical & Computer Engineering

**Computer Science Award for Outstanding Research**

USC Viterbi School of Engineering

## 2022 US Granted Patents

**A.R. Jaiswal, B.C. Paul, S.R. Soss**

*Non-volatile transistor embedded static random access memory (SRAM) cell* (US Patent 11,475,941)

**A.R. Jaiswal, A.P. Jacob**

*Array of integrated pixel and memory cells for deep in-sensor, in-memory computing* (US Patent 11,468,146)

**A.R. Jaiswal, A.P. Jacob, Y. Bian, D.C. Pritchard**

*Image sensor incorporating an array of optically switchable magnetic tunnel junctions* (US Patent 11,226,231)

**A.R. Jaiswal, A.P. Jacob, Y. Bian, M. Rakowski**

*Optical neuro-mimetic devices* (US Patent 11,537,866)

**K. Karuppiah, G. Sankaran**

*In-band-telemetry-based path MTU size determination system* (US Patent 11,405,325)

## NEW RESEARCH DIRECTORS, TEAM LEADERS, AND LEADS

At ISI, **Research Directors** head a sizable scientific research group and develop, plan, and direct the research and funding strategy for that group while also contributing to the current and future success of the institute. **Research Team Leaders** lead sizable research groups and are instrumental in seeking out and recruiting other scientists and researchers to ensure ISI's technical integrity and its reputation as a leader. **Research Leads** continuously make noteworthy contributions to their research field and to ISI's technological advancement.

### RESEARCH DIRECTOR



**John Paul Walters**  
*Computational Systems  
and Technology Division*

### RESEARCH TEAM LEADERS



**Itay Hen**  
*Computational Systems  
and Technology Division*



**Alefiya Hussain**  
*Networking and  
Cybersecurity Division*


### RESEARCH LEADS



**Jeremy Abramson**  
*Networking and  
Cybersecurity Division*



**Muhao Chen**  
*Artificial Intelligence  
Division*



**Michael Collins**  
*Networking and  
Cybersecurity Division*



**Luis Garcia**  
*Networking and  
Cybersecurity Division*



**Wesley Hardaker**  
*Networking and  
Cybersecurity Division*



**Abigail Horn**  
*Executive Office*



**Filip Ilievski**  
*Artificial Intelligence  
Division*



**Brian Kocoloski**  
*Networking and  
Cybersecurity Division*



**Xuezhe Ma**  
*Artificial Intelligence  
Division*



**Mohammad Rostami**  
*Artificial Intelligence  
Division*



**Leonidas Spinoulas**  
*Artificial Intelligence  
Division*





**Emilio Ferrara** was promoted to **Full Professor**. He holds a joint appointment in the Annenberg School for Communication and Journalism and the Department of Computer Science, as well as an appointment in Preventive Medicine.



**Luis Garcia**, in the Networking and Cybersecurity Division, is now **Research Assistant Professor** in the Computer Science department.




**Abigail Horn** received an appointment as **Research Assistant Professor** in the Industrial and Systems Engineering department.



**Filip Ilieveski**, in the Artificial Intelligence Division, is now **Research Assistant Professor** in the Computer Science department.



**Jonathan May**, in the Artificial Intelligence Division, was promoted to **Research Associate Professor** of Computer Science.



## 2022 PhD GRADUATES

ISI is proud to congratulate these stellar graduates who earned their PhD degrees or successfully defended their PhD dissertations in 2022.

### **Ayush Jaiswal**

Dissertation: *Invariant Representation Learning*

Current position: Applied Computer Scientist, Amazon

### **Ekraam Sabir**

Dissertation: *Detection of Semantically Repurposed Natural and Biomedical Images*

Current position: Applied Computer Scientist, Amazon

**Wael AbdAlmageed, Advisor**

### **Tu Mai Anh Do**

Dissertation: *Optimizing Execution of In Situ Workflows*

Current position: Machine Learning Engineer, MoMo

**Ewa Deelman, Advisor**

### **Akira Matsui**

Dissertation: *Mining and Modeling Temporal Structures of Human Behavior in Digital Platforms*

Current position: Assistant Professor, Yokohama National University

### **Karishma Sharma**

Dissertation: *Diffusion Network Inference and Analysis for Disinformation Mitigation*

Current position: Applied Scientist, Amazon Alexa AI

### **Shen Yan**

Dissertation: *Fair Machine Learning for Human Behavior Understanding*

Current position: Research Scientist, Meta, Facebook Core Data Science

**Emilio Ferrara, Advisor**

### **Ninareh Mahrabi**

Dissertation: *Responsible Artificial Intelligence for a Complex World*

Current position: Postdoctoral Scholar, Google

(Fred Morstatter, Co-advisor)

**Aram Galstyan, Advisor**

### **Robert Schuler**

Dissertation: *Schema Evolution for Scientific Asset Management*

Current position: Senior Computer Scientist and Research Lead, ISI

**Carl Kesselman, Advisor**

### **Minh Pham**

Dissertation: *Robust and Proactive Error Detection and Correction in Tables*

Current position: Software Engineer, Google

**Craig Knoblock, Advisor**

### **Nada Aldarrab**

Dissertation: *Automatic Decipherment of Historical Manuscripts*

Current position: Assistant Professor, King Abdulaziz University

### **Thamme Gowda**

Dissertation: *The Inevitable Problem of Rare Phenomena Learning in Machine Translation*

Current position: Research Scientist, Microsoft Translator Team

**Jonathan May, Advisor**

### **Rajat Tandon**

Dissertation: *Protecting Online Services from Sophisticated DDoS Attacks*

Current position: Software Engineer, Juniper Networks

**Jelena Mirkovic, Advisor**

### **Victor Ardulov**

Dissertation: *Modeling and Regulating Human Interaction with Control Affine Dynamical Systems*

Current position: Applied Scientist, Amazon

### **Nikolaos Flemotomos**

Dissertation: *Extracting and Using Speaker Role Information in Speech Processing Applications*

Current position: Machine Learning Researcher/Engineer, Apple

### **Raghuv eer Peri**

Dissertation: *Neural Representation Techniques for Robust and Fair Speaker Recognition*

Current position: Applied Scientist, Amazon Web Services AI Lab

**Shri Narayanan, Advisor**

### **Serban Stan**

Dissertation: *Unsupervised Domain Adaptation with Private Data*

Current position: Research Scientist, Facebook

**Mohammad Rostami, Advisor**

### **Sami Abu-El-Hajja**

Dissertation: *Fast and Label-Efficient Graph Representation Learning*

Current position: Senior Research Scientist, Google Research

### **Rob Brekelmans**

Dissertation: *Information Geometry of Annealing Paths for Inference and Estimation*

Current position: Postdoctoral Fellow, Vector Institute

### **Hannes Leipold**

Dissertation: *Imposing Classical Symmetries on Quantum Operators with Applications to Optimization*

Current position: Graduate Research Assistant, ISI

Federico Spedalieri, Co-advisor

**Greg Ver Steeg, Advisor**

## Visiting Researchers

**Tommaso Fornaciari**, Bocconi University  
*Deception detection models*

**Francisco Pierri**, Politecnico di Milano  
*Moderation mechanisms of online platforms*

**Emilio Ferrara**, *Host*

**Yagi Tomonori**, Japan Patent Office

*Creating signatures for orbital debris using various sensors*

**David Barnhart**, *Host*

## Postdoctoral Scholars

**Emre Akaturk** | *Quantum Monte Carlo simulations and quantum computing*

**Itay Hen**, *Supervisor*

**Emmanuel Johnson** | *Utilizing artificial intelligence methods to build personalized learning systems*

**Yolanda Gil**, *Supervisor*

**Luca Luceri** | *Machine learning, data science, and network science to investigate, detect, and mitigate online harms in socio-technical systems*

**Emilio Ferrara**, *Supervisor*

**Mitch Paul Mithun** | *Quantum natural language processing, machine learning models for phishing email detection*

**Marjorie Freedman**, *Supervisor*

**Tahereh Rezaei** | *Experimental quantum measurement for enabling new applications in quantum communications*

**Jonathan Habif**, *Supervisor*

**Rahul Rughani** | *Spacecraft rendezvous and docking*

**David Barnhart**, *Supervisor*

**Sugheet Sunder** | *Silicon photonics devices, fiber-to-chip I/Os, optical logic gates*

**Ajey Jacob**, *Supervisor*



Some new PhDs in the Artificial Intelligence Division celebrate

## PHD STUDENTS

ISI is proud to recognize this impressive cohort of doctoral students and their advisors. The PhD students here include those who are supported at ISI and students whose advisors are associated with ISI.

During the summer, many ISI PhD students interned at leading tech companies, including Apple, IBM, Bloomberg, Amazon Alexa AI, and the Aerospace Corporation.



**Jiaxin Cheng** | *Computer vision, style transfer, image synthesis, adversarial detection*

Prem Natarajan (Alexa AI), Co-advisor

**Jiazhi Li** | *Debiasing face recognition*

**Soumyaroop Nandi** | *Computer vision and deep learning in image forensics*

Prem Natarajan (Alexa AI), Co-advisor

**Hae Jin (Hayley) Song** | *Generative modeling and representation learning*

**Mulin Tian** | *Deepfake detection*

**Hanchen Xie** | *Learning intuitive physics*

**Jiageng Zhu** | *Learning causal structures*

**Wael AbdAlmageed, Advisor**

---

**Dimitrios Stripelis** | *Federated learning*

**Jose Luis Ambite, Advisor**

---

**David Bacher** | *Autonomous science targeting and divert guidance during entry, descent, and landing phases*

**Cristobal Garrido** | *Reducing error uncertainties in rendezvous proximity operations for small satellites*

**Jonathan Messer** | *Solar pressure to reduce tumble rate on small satellites*

**Adarsh Rajguru** | *Small satellite rendezvous and docking optimization*

**David Barnhart, Advisor**

**James Y. Huang** | *Natural language processing*

**Tenghao Huang** | *Multi-hop reasoning for natural language processing*

**Bangzheng Li** | *Indirectly supervised natural language processing*

**Qin Liu** | *Learnability and denoising*

**Ehsan Qasemi** | *Natural language processing, common sense reasoning, neuro-symbolic AI*

**Fei Wang** | *Robustness and accountability in natural processing*

**Peifeng Wang** | *Common sense reasoning*

**Nan Xu** | *Accountability and generalizable information extraction*

**Wenxuan Zhou** | *Natural language processing*

**Muhao Chen, Advisor**

---

**Haonan Wang** | *Edge computing*

**Stephen Crago, Advisor**

---

**Patrycja Krawczuk** | *Machine learning for workflow management systems*

**George Papadimitriou** | *Distributed computing and scientific workflows*

**Ewa Deelman, Advisor**

**Charles (Duke) Bickham** | *AI fairness in healthcare and mental health*

**Alexander Bisberg** | *Pro-social behavior and collaboration in multiplayer online games*

**Herbert Chang** | *Improving information diversity on social media*

**Emily Chen** | *Computational social science; human behavior, influence, and information flow in social platforms*

**Eun Cheol Choi** | *Reasoning on large language models*

**Bijean Ghafouri** | *Computational survey of political knowledge*

**Julie Jiang** | *Computational social science*

**Hsien-Te Kao** | *Algorithmic fairness*

**Alex Spangher** | *Computational journalism*

**Yilei Zeng** | *Modeling human sequential decision-making in online platforms and gaming*

**Meiqing Zhang** | *Dimensionality of political cleavages*

**Emilio Ferrara, Advisor**

**Coleman, Tainã** | *Scientific workflow structure detection, benchmarking, and scheduling*

**Rafael Ferreira da Silva**  
(Oak Ridge National Laboratory), *Advisor*

**Sarik Ghazarian Ghalemaleki** | *Natural language processing, automatic evaluation of open-domain natural language generation*

**Aram Galstyan, Advisor**

**Haley Weinstein** | *Quantum optical communication*

**Jonathan Habif, Advisor**

**Guillermo Baltra Elorriaga** | *Internet reliability*

**Asma Enayet** | *Internet outage detection from passive data sources*

**Basileal Yoseph Imana** | *Assessing the algorithmic fairness of social media platforms*

**Aleksandra Korolova** (Princeton), *Co-advisor*

**ASM Rizvi** | *Networking, security, and Internet measurement*

**Xiao Song** | *Internet measurement*

**Kicho Yu** | *Network security*

**John Heidemann, Advisor**

**Babakhani, Arman** | *Quantum Monte Carlo simulations*

**Itay Hen, Advisor**

**Zhivar Sourati** | *Common sense reasoning with natural language processing and knowledge graphs*

**Jiarui Zhang** | *Common sense reasoning with natural language processing and knowledge graphs*

**Filip Ilievski, Advisor**

**Md Abdullah-Al Kaiser** | *Analog and mixed-signal integrated circuit design*

**Zihan Yin** | *In-pixel in-memory computing circuit design*

**Akhilesh Jaiswal, Advisor**

**Ke Shen** | *Machine common sense*

**Mayank Kejriwal, Advisor**

**Fandel Lin** | *Geographic information science, data mining, machine learning, data integration, and urban computing*

**Basel Shbita** | *Automatic semantic spatio-temporal interpretation of digitized historical maps*

**Binh Vu** | *Automatic data modeling*

**Craig Knoblock, Advisor**

**Nathan Bartley** | *Socially embedded algorithms*

**Rebecca Dorn** | *Trans-inclusive research frameworks and gender bias in data*

**Dan Feldman** | *Socially embedded algorithms*

**Fiona Guo** | *Computational social sciences*

**Zihao He** | *Computational social science and natural language processing*

**Ashwin Shreyas Mohan Rao** | *Computational social science and social networks*

**Negar Mokhberian** | *Estimation of dataset quality in subjective tasks and evaluating how human biases affect model decisions*

**Cassandra Rusti** | *Fairness and bias*

**Kristina Lerman, Advisor**

**Jiao Sun** | *Trustworthy text generation*

**Nuan Wen** | *Natural language understanding from the perspective of pragmatics*

**Xuezhe Ma, Advisor**

**Hyundong (Justin) Cho** | *Natural language processing in dialogue systems*

**Katy Felkner** | *Fairness and bias in large language models*

**Mozhdeh Gheini** | *Transfer learning for machine translation*

**Meryem M'Hamdi** | *Cross-lingual transfer learning*

**Jonathan May, Advisor**

**Pithayuth Charnsethikul** | *Human factors and cybersecurity*

**Nicolaas Weideman** | *Binary program analysis for vulnerability detection and mitigation*

**Wei-Cheng Wu** | *Computer security, including software testing, binary analysis, and IoT firmware security*

**Christophe Hauser, Luis Garcia, Co-advisors**

**Jelena Mirkovic, Advisor**

## PHD STUDENTS (CONTINUED)

**Bahareh Harandizadeh** | *Causal modeling of online to offline dynamics*

**Yuzhong Huang** | *Computer vision and 3D reconstruction*

**Abel Salinas** | *Natural language processing*

**Zhuoyu Shi** | *Computational social science, social network analysis*

**Fred Morstatter, Advisor**

---

**Sabyasachee Baruah** | *Computational media intelligence*

**Zhuohao Chen** | *Speech and language processing*

**Georgios Chochlakis** | *Scalable and flexible emotion recognition on social media*

**Tiantian Feng** | *Trustworthy machine intelligence*

**Amrutha Nadarajan** | *Machine learning, wearable sensing, audio processing*

**Shrikanth Narayanan, Advisor**

---

**I-Hung Hsu** | *Information extraction and model generalizability in natural language processing*

**Nanyun Peng (UCLA), Advisor**

**Prem Natarajan (Alexa AI), Advisor**

---

**Kian Ahrabian** | *Machine learning on graphs*

**Pegah Jandaghimeibodi** | *Dialog systems, conversational recommenders, transfer learning*

**Dong-Ho Lee** | *Low-resource natural language processing, explanation-based learning, knowledge graphs*

**Kexuan Sun** | *Table understanding, knowledge graphs*

**Avijit Thawani** | *Natural language processing, tokenization*

**Pei Zhou** | *Common ground reasoning for conversational AI, natural language processing, machine learning*

Xiang Ren, Co-advisor

**Jay Pujara, Advisor**

---

**Sima Arasteh** | *Software security and binary analysis*

**Mukund Raghothoman (USC), Advisor**

---

**Weizhao Jin** | *Secure computation and privacy-preserving solutions for next-gen networks and distributed systems*

**Yixiang Yao** | *Applying privacy-enhancing techniques to AI applications*

**Srivatsan Ravi, Advisor**

---

**Arakelyan, Shushan** | *Machine learning models for code*

**Jake Bremerman** | *Natural language processing*

**Chan, Aaron** | *Model explainability*

**Woojeong Jin** | *Natural language processing, vision-language models, common sense reasoning*

**Brihi Joshi** | *Interpretable natural language processing*

**Huihan Li** | *Training and evaluating language models*

**Yuchen (Bill) Lin** | *Commonsense reasoning, federated learning*

**Sahana Ramnath** | *Natural language processing, deep learning, reinforcement learning*

**Soumya Sanyal** | *Logical reasoning in language models*

**Jun Yan** | *Trustworthy natural language processing*

**Qinyuan Ye** | *Natural language processing and few-shot learning*

**Xiang Ren, Advisor**

---

**Ang Li** | *Data mining*

**Yi Zheng** | *Artificial intelligence and networking*

**Satish Kumar Thittamarahalli, Advisor**

---

**Tejas Srinivasan** | *Multimodal continual learning*

**Jesse Thomason (USC), Advisor**

---

**Umang Gupta** | *Fairness and privacy in machine learning*

**Hrayr Harutyunyan** | *Information stored in neural network weights or activations and its connections to generalization, memorization, stability, and learning dynamics*

Aram Galstyan, Co-advisor

**Neal Lawton** | *Machine learning, optimization*

**Elan Markowitz** | *Scalable graph learning and applications to knowledge graphs*

Aram Galstyan, Co-advisor

**Myrl Marmarelis** | *Causal inference, high-dimensional statistics, and time series*

Aram Galstyan, Co-advisor

**Greg Ver Steeg, Advisor**

---

## Rising Star MS Internship Program

ISI offers Rising Star Internships to students applying to the Viterbi School of Engineering Master of Science program in computer science, data science, electrical and computer engineering, and related disciplines. The interns collaborate on an ongoing project with a research team—and work directly with an individual mentor throughout the course of the master’s degree. Currently, one intern is working with researchers in the Informatics Systems Research Division on data management for the World in a Cell transmedia project with computer scientists, biologists, and film school researchers; another intern is developing an automated shell script parser to perform safe de-obfuscation of malicious shell scripts. Past interns worked in analyzing malware samples and knowledge graph profiles. Scan QR code for more information.



## Summer Internship Program

ISI welcomes PhD students, master’s degree students, and undergraduates every summer as paid interns to work with senior research leaders and their teams. The internships are available at all three ISI locations, and interns may work on-site or remotely with their ISI mentors and research teams.

Interns work in all ISI research divisions in data science, hardware security, machine learning, natural language processing, quantum information science, and other areas. In addition to research, interns participate in seminars, reading groups, and social activities. Scan QR code for more information.



## Research Experiences for Undergraduates

The NSF Research Experiences for Undergraduates (REU) program offers an intellectually and socially stimulating paid research experience to undergrads from a broad range of colleges and universities at ISI headquarters in Marina del Rey, CA. In addition to the stipend, students are provided with campus housing at USC and free shuttle service between campus and ISI. Each student concentrates on a specific research project and works closely with ISI research faculty and other ISI researchers.

The broad research theme is *Safe, Usable, Fair, and Reliable Internet*: communication and cybersecurity, scientific experimentation and knowledge capture, and social data science. The program is led by Jelena Mirkovic, an ISI cybersecurity researcher and project leader. Over the summer, the undergrads work with ISI researchers, individually and in teams, and attend seminars and social events. At the conclusion of the program, the students present their research as poster sessions at the ISI REU symposium to their ISI mentors, researchers, and students. Scan QR code for more information.



## Computing Architectures and Systems

### An Automated Cryo-EM Computational Environment on the HPC System using Pegasus WMS

T. Osinski, M. Rynge, J. K. Hong, K. Vahi, R. Chu, C. Sul, E. Deelman, and B.D. Kim  
*2022 IEEE/ACM Workshop on Workflows in Support of Large-Scale Science (WORKS)*

### Bitstream Assurance Checking Engine for Undocumented Functionality

A. Schmidt, J. Wilford, B. Reynwar, T. Sung, and M. French  
*Government Microcircuit Applications and Critical Technology Conference (GOMACTech)*

### FaceBase: A Community-Driven Hub for Data-Intensive Research

R. Schuler, A. Bugacov, J.G. Hacia, T. Ho, J. Iwata, L. Pearlman, B. Samuels, C. Williams, Z. Zhao, C. Kesselman, and Y. Chai  
*Journal of Dental Research*

### GOAL: Supporting General and Dynamic Adaptation in Computing Systems

A. Pervaiz, Y.H. Yang, A. Duracz, F. Bartha, R. Sai, C. Imes, R. Cartwright, K. Palem, S. Lu, and H. Hoffmann  
*Proceedings of the 2022 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*

### MusMorph, a Database of Standardized Mouse Morphology Data for Morphometric Meta-Analyses

J. Devine, M. Vidal-García, W. Liu, A. Neves, L. Lo Vercio, R. Green, H. Richbourg, M. Marchini, C. Unger, A. Nickle, B. Radford, N. Young, P. Gonzalez, and R. Schuler  
*Scientific Data*

### P2M-DeTrack: Processing-in-Pixel-in-Memory for Energy-efficient and Real-Time Multi-Object Detection and Tracking

G. Datta, S. Kundu, Z. Yin, J. Mathai, Z. Liu, Z. Wang, M. Tian, S. Lu, R. T. Lakkireddy, Andrew G. Schmidt, W. AbdAlmageed, A. P. Jacob, A. R Jaiswal, P. A. Beerel  
*IFIP IEEE 30th International Conference on Very Large Scale Integration (VLSI SoC)*

### StereoBit on the SpaceCube Mini

J. Carr, C. Wilson, D. Wu, M. French, M. Paolieri, H. Madani, and M. Kelly  
*IEEE International Geoscience and Remote Sensing Symposium*

### Towards Full-Stack Acceleration for Fully Homomorphic Encryption

N. Zhang, H. Gamil, P. Brinich, B. Reynwar, A. Badawi, N. Neda, D. Soni, K. Canida et al.  
*IEEE High Performance Extreme Computing Conference (HPEC)*

### Untangling IP Protection via Learning and Structure

D. Chen, X. Zhou, S. Chowdhury, P. Beerel, P. Nuzzo, and M. French  
*Conference, Government Microcircuit Applications and Critical Technology Conference (GOMACTech)*

## Cybersecurity

### A Tutorial on Security and Privacy Challenges in CPS

S.M. Dibaji, A. Hussain, and H. Ishii  
*Security and Resilience of Control Systems*

### AMON-SENS: Scalable and Accurate Detection of Volumetric DDoS Attacks at ISPs

R. Tandon, P. Charnsethikul, M. Kallitsis, and J. Mirkovic  
*Globecom*

### Arbiter: Bridging the Static and Dynamic Divide in Vulnerability Discovery on Binary Programs

J. Vadayath, M. Eckert, K. Zeng, N. Weideman et al.  
*USENIX Security*

### AutoCPS: Control Software Dataset Generation for Semantic Reverse Engineering

H. Wang, C. Hauser, and L. Garcia  
*2022 IEEE Security and Privacy Workshops (SPW)*

### Conflict-Based Search for the Virtual Network Embedding Problem

Y. Zheng, S. Ravi, E. Kline, S. Koenig, and T.K. Satish Kumar  
*Proceedings of the Thirty-Second International Conference on Automated Planning and Scheduling (ICAPS 2022)*

### Decentralized Privacy-Preserving Path Validation for Multi-Slicing-Authority 5G Networks

W. Jin, S. Ravi, and E. Kline  
*IEEE Wireless Communications and Networking Conference*

### Dynamic Graph Reduction Optimization Technique for Interdiction Games

J. Blythe and A. Tregubov  
*AAMAS Workshop on Optimization and Learning in Multiagent Systems*

### Generating Representative Video Teleconferencing Traffic

D. DeAngelis, A. Hussain, B. Kocoloski, C. Ardi, and S. Schwab  
*Proceedings of the 15th Workshop on Cyber Security Experimentation and Test*

### Harm-DoS: Hash Algorithm Replacement for Mitigating Denial-of-Service Vulnerabilities in Binary Executables

N. Weideman, H. Wang, T. Kann, S. Zahabizadeh, W. Wu, R. Tandon, J. Mirkovic, and C. Hauser  
*RAID*

### I Know What You Did on Venmo: Discovering Privacy Leaks In Mobile Social Payments

R. Tandon, P. Charnsethikul, I. Arora, D. Murthy, and J. Mirkovic  
*Privacy Enhancing Technologies Symposium (PETS)*

### Large-Scale Agent-Based Simulations of Online Social Networks

G. Muric, A. Tregubov, J. Blythe, A. Abeliuk, D. Choudhary, K. Lerman, and E. Ferrara  
*Autonomous Agents and Multi-Agent Systems*



**Modeling Human-Cyber Interactions in Safety-Critical Cyber-Physical/Industrial Control Systems**

S. Ngo, D. DeAngelis, and L. Garcia  
*2022 IEEE 19th International Conference on Mobile Ad Hoc and Smart Systems (MASS)*

**NS3: Neuro-Symbolic Semantic Code Search**

S. Arakelyan, A. Hakhverdyan, M. Allamanis, C. Hauser, L. Garcia, and X. Ren  
*Advances in Neural Information Processing Systems*

**Polymorphic Malware Behavior through Network**

**Trace Analysis**  
 X. Deng and J. Mirkovic  
*Proceedings of 14th International Conference on COMmunication Systems & NETworks (COMSNETS)*

**Reverse Engineering and Retrofitting Robotic Aerial Vehicle Control Firmware using Dispatch**

T. Kim, A. Ding, S. Etigowni, P. Sun, J. Chen, L. Garcia, S. Zonouz, D. Xu, and D. Tian  
*Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services*

**Safety in the Emerging Holodeck Applications**

S. Ghandeharizadeh and L. Garcia  
*Proceedings of the 1st Workshop on Novel Challenges of Safety, Security and Privacy in Extended Reality*

**Samba: Identifying Inappropriate Videos for Young Children on YouTube**

L. Binh, R. Tandon, C. Oinar, J. Liu, U. Durairaj, J. Guo, S. Zahabizadeh, S. Ilango, J. Tang, F. Morstatter, S. Woo, and J. Mirkovic  
*31st ACM International Conference on Information & Knowledge Management (CKIM)*

**Secure Publish-Process-Subscribe System for Dispersed Computing**

W. Jin, B. Krishnamachari, M. Naveed, S. Ravi, E. Sanou, and K. Wright  
*41st International Symposium on Reliable Distributed Systems (SRDS 2022)*

**The DARPA SEARCHLIGHT Dataset of Application Network Traffic**

C. Ardi, C. Aubry, B. Kocoloski, D. DeAngelis, A. Hussain, M. Troglia, and S. Schwab  
*Proceedings of the 15th Workshop on Cyber Security Experimentation and Test*

**The FastMap Pipeline for Facility Location Problems**

O. Thakoor, A. Li, S. Koenig, S. Ravi, E. Kline, and T. K. S. Kumar  
*PRIMA 2022: Principles and Practice of Multi-Agent Systems - 24th International Conference*

**The Limits of Helping in Non-volatile Memory Data Structures**

O. Ben-Baruch and S. Ravi  
*Stabilization, Safety, and Security of Distributed Systems - 24th International Symposium (SSS 2022)*

**Tinyodom: Hardware-Aware Efficient Neural Inertial Navigation**

S.S. Saha, S.S. Sandha, L. Garcia, and M. Srivastava  
*Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*

**Toward Findable, Accessible, Interoperable, and Reusable Cybersecurity Artifacts**

D. Balenson, T. Benzel, E. Eide, D. Emmerich, D. Johnson, J. Mirkovic, and L. Tinnel  
*Cyber Security Experimentation and Test (CSET)*

**Xatu: Boosting Existing DDoS Detection Systems Using Auxiliary Signals**

Z. Xu, S. Ramanathan, A. Ruth, J. Mirkovic, and M. Yu  
*CoNext*

**PREP-UC: A Practical Replicated Persistent Universal Construction**

G. C. Coccimiglio, T. A. Brown, and S. Ravi  
*SPAA 22: 34th ACM Symposium on Parallelism in Algorithms and Architectures*

**Knowledge Technologies**

**Alzheimer’s Disease Detection with a 3D Convolutional Neural Network using Gray Matter Maps from T1-weighted Brain MRI**

N. Dhinagar, S. Thomopoulos, C. Owens-Walton, D. Stripelis, J. Ambite, G. ver Steeg, and P. Thompson  
*Alzheimer’s & Dementia*

**An Empirical Investigation of Commonsense Self-Supervision with Knowledge Graphs**

J. Zhang, F. Ilievski, Filip, K. Ma, J. Francis, and A. Oltramari  
*AKBC 2022*

**Augmenting Knowledge Graphs for Better Link Prediction**

J. Wang, F. Ilievski, P. Szekely, and K. Yao  
*IJCAI-ECAI*

**Automated Generation of Control Concepts Annotation Rules Using Inductive Logic Programming**

B. Shbita and A. Moitra  
*Functional and Logic Programming*

**Building Spatio-Temporal Knowledge Graphs from Vectorized Topographic Historical Maps**

B. Shbita, C. Knoblock, W. Duan, Y. Chiang, J. Uhl, and S. Leyk  
*Semantic Web*

### Coalescing Global and Local Information for Procedural Text Understanding

K. Ma, F. Ilievski, J. Francis, E. Nyberg, and A. Oltramari  
*COLING 2022*

### Dataset for Studying Gender Disparity in English Literary Texts

A. Nagaraj and M. Kejriwal  
*Elsevier Data in Brief*

### Decision Making in Monopoly Using a Hybrid Deep Reinforcement Learning Approach

T. Bonjour, M. Haliem, A.O. Alsalem, S. Thomas, H. Li, V. Aggarwal, M. Kejriwal, and B.K. Bhargava  
*IEEE Transactions on Emerging Topics in Computational Intelligence*

### Designing a Strong Test for Measuring True Common-Sense Reasoning

M. Kejriwal, H. Santos, A.M. Mulvehill, and D.L. McGuinness  
*Nature Machine Intelligence*

### Does Wikidata Support Analogical Reasoning?

F. Ilievski, J. Pujara, and K. Shenoy  
*Iberoamerican Knowledge Graphs and Semantic Web Conference (KGSWC)*

### Enabling Scientific Reproducibility through FAIR Data Management: An Ontology-Driven Deep Learning Approach in the NeuroBridge Project

X. Wang, Y. Wang, and J. Ambite et al.  
*American Medical Informatics Association (AMIA) 2022 Annual Symposium*

### Enriching Wikidata with Linked Open Data

B. Zhang, F. Ilievski, and P. Szekely  
*Wikidata-22 Workshop*

### Evaluation of Transfer Learning Methods for Detecting Alzheimer's Disease with Brain MRI

N. Dhinagar, S. Thomopoulos, P. Rajagopalan, D. Stripelis et al.  
*18th International Symposium on Medical Information Processing and Analysis (SIPAIM)*

### Federated Deep Learning for Detecting Alzheimer's Disease in Multi-Cohort Brain MRI

D. Stripelis, N. Dhinagar, R. Sanchez Romero et al.  
*Alzheimer's Association International Conference*

### Federated Progressive Sparsification (Purge-Merge-Tune)

D. Stripelis, U. Gupta, G. Ver Steeg, and J. Ambite  
*Workshop on Federated Learning: Recent Advances and New Challenges (in conjunction with NeurIPS 2022)*

### Generalizable Neuro-symbolic Systems for Commonsense Question Answering

A. Oltramari, J. Francis, F. Ilievski, K. Ma, and R. Mirzaee  
*Neuro-Symbolic Artificial Intelligence: The State of the Art*

### Identifying Surprising Facts in Wikidata

N. Klein, F. Ilievski, H. Freedman, and P. Szekely  
*Wikidata-22 Workshop*

### KGTK: Tools for Creating and Exploiting Large Knowledge Graphs (tutorial)

R. Tommasini, B. Riccaran, S. Roy, X. Wang, H. Wang, H. Ji, J. Han, P. Nakov, G. Da San Martino, F. Alam, M. Schedl, F. Ilievski et al.  
*Companion Proceedings of the Web Conference 2022 and Knowledge Graph Conference 2022 (KGC)*

### Knowledge Graphs for Social Good: An Entity-Centric Search Engine for the Human Trafficking Domain

M. Kejriwal and P. Szekely  
*IEEE Transactions on Big Data*

### Knowledge Graphs: A Practical Review of the Research Landscape

Mayank Kejriwal  
*Information*

### PINTO: Faithful Language Reasoning Using Prompt-Generated Rationales

P. Wang, A. Chan, F. Ilievski, M. Chen, and X. Ren  
*TSRML 2022 Workshop*

### Predicted Gene Expression in Ancestrally Diverse Populations Leads to Discovery of Susceptibility Loci for Lifestyle and Cardiometabolic Traits

H. Highland, G. Wojcik, M. Graff, K. Nishimura, J. Ambite et al.  
*American Journal of Human Genetics*

### SAND: A Tool for Creating Semantic Descriptions of Tabular Sources

B. Vu and C. Knoblock  
*The Semantic Web: ESWC 2022 Satellite Events*

### Semi-Synchronous Federated Learning for Energy-Efficient Training and Accelerated Convergence in Cross-Silo Settings

D. Stripelis, P.M. Thompson, and J. Ambite  
*ACM Transactions on Intelligent Systems and Technology, Special Issue on Federated Learning: Algorithms, Systems, and Applications*

### Shared Genetics Drive Mate Selection

K. Burghardt, T.E. Arpawong, and J. Ambite  
*Annual Meeting of the American Society of Human Genetics*

### Towards Sparsified Federated Neuroimaging Models via Weight Pruning

D. Stripelis, U. Gupta, N. Dhinagar, G. Ver Steeg, P. Thompson, and J. Ambite  
*Distributed, Collaborative, and Federated Learning, and Affordable AI and Healthcare for Resource Diverse Global Health - Third MICCAI Workshop, DeCaF 2022, and Second MICCAI Workshop, FAIR 2022*

**Towards the Automated Large-Scale Reconstruction of Past Road Networks from Historical Maps**

J.H. Uhl, S. Leyk, Y. Chiang, and C. Knoblock  
*Computers, Environment and Urban Systems*

**Understanding Narratives through Dimensions of Analogy**

T. Nagarajah, F. Ilievski, and J. Pujara  
*Workshop on Qualitative Reasoning (QR)*

**Utilizing Background Knowledge for Robust Reasoning over Traffic Situations**

J. Zhang, F. Ilievski, A. Kollaa, J. Francis, K. Ma, and A. Oltramari  
*Workshop on Knowledge Augmented Methods for NLP*

**Machine Learning and Data Science**

**An AI Approach to Integrating Climate, Hydrology, and Agriculture Models**

B. Berhanu, E. Bisrat, Y. Gil, D. Khider, M. Osorio, V. Ratnakar, and H. Vargas  
*Proceedings of the First International Workshop on Social Impact of AI for Africa (SIAIA), held at the 36th Annual Conference of the Association for the Advancement of Artificial Intelligence (AAAI-22)*

**Assessing Scientific Research Papers with Knowledge Graphs**

K. Sun, Z. Qiu, A. Salinas, Y. Huang, D. Lee, D. Benjamin, F. Morstatter, X. Ren, K. Lerman, J. Pujara  
*Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*

**Automating Detection of Papilledema in Pediatric Fundus Images with Explainable Machine Learning**

K. Kleanthis, M. Rostami, M. Chang, S. Narayanan  
*2022 IEEE International Conference on Image Processing (ICIP)*

**Charting the Information and Misinformation Landscape to Characterize Misinfodemics on Social Media: COVID-19 Infodemiology Study at a Planetary Scale**

E. Chen, J. Jiang, H. Chang, G. Muric, and E. Ferrara  
*JMIR Infodemiology*

**CLiMB: A Continual Learning Benchmark for Vision-and-Language Tasks**

T. Srinivasan, T. Chang, L. Alva, G. Chochlakis, M. Rostami, Mohammad, and J. Thomason  
*Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*

**Domain Adaptation for the Segmentation of Confidential Medical Images**

S. Stan and M. Rostami  
*British Machine Vision Conference*

**Embedding Physical Flow Functions into Deep Learning Predictive Models for Improved Production Forecasting**

S. Razak, J. Cornelio, Y. Cho, H. Liu, R. Vaidya, and B. Jafarpour  
*Unconventional Resources Technology Conference*

**Estimating Numbers Without Regression**

A. Thawani, J. Pujara, and A. Kalyan  
*NeurIPS Workshop on MathAI*

**Failure Modes of Domain Generalization Algorithms**

T. Galstyan, H. Harutyunyan, H. Khachatrian, G. ver Steeg, and A. Galstyan  
*Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*

**FETA: A Benchmark for Few-Sample Task Transfer in Open-Domain Dialogue**

A. Albalak, Y. Tuan, P. Jandaghi, C. Pryor, L. Yoffe, D. Ramachandran, L. Getoor, J. Pujara, and W. Wang  
*Conference on Empirical Methods in Natural Language Processing*

**Formal Limitations of Sample-Wise Information-Theoretic Generalization Bounds**

H. Harutyunyan, G. ver Steeg, and A. Galstyan  
*2022 IEEE Information Theory Workshop (ITW)*

**Gendered Citation Patterns Among the Scientific Elite**

K. Lerman, Y. Yu, F. Morstatter, and J. Pujara  
*Proceedings of the National Academy of Sciences*

**Good Examples Make a Faster Learner: Simple Demonstration-Based Learning for Low-Resource NER**

D. Lee, A. Kadakia, K. Tan, M. Agarwal, X. Feng, T. Shibuya, R. Mitani, T. Sekiya, J. Pujara, and X. Ren  
*Association for Computational Linguistics*

**Identifying Geopolitical Event Precursors using Attention-Based LSTMs**

K. Hossain, H. Harutyunyan, Y. Ning, B. Kennedy, N. Ramakrishnan, and A. Galstyan  
*Frontiers in Artificial Intelligence*

**Increasing Model Generalizability for Unsupervised Visual Domain Adaptation**

M. Rostami  
*Conference on Lifelong Learning Agents*

**Inferring Topological Transitions in Pattern-Forming Processes with Self-Supervised Learning**

M. Abram, K. Burghardt, G. ver Steeg, A. Galstyan, and R. Dingreville  
*npj Computational Materials*

**Learning Fairer Interventions**

Y. He, K. Burghardt, S. Guo, and K. Lerman  
*Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*

**Leveraging Change Point Detection to Discover Natural Experiments in Data**

Y. He, K. Burghardt, and K. Lerman  
*EPJ Data Science*

### **Leveraging Visual Knowledge in Language Tasks: An Empirical Study on Intermediate Pre-Training for Cross-Modal Knowledge Transfer**

W. Jin, D. Lee, C. Zhu, J. Pujara, and X. Ren

*Association for Computational Linguistics*

### **Machine Learning on Graphs: A Model and Comprehensive Taxonomy**

I. Chami, S. Abu-El-Haija, B. Perozzi, C. Re, and K. Murphy

*Journal of Machine Learning Research*

### **Noise Audits Improve Moral Foundation Classification**

N. Mokhberian, F. Hopp, B. Harandizadeh, F. Morstatter, and K. Lerman

*Proceedings of the 2023 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*

### **Partisan Asymmetries in Exposure to Misinformation**

A. Rao, F. Morstatter, and K. Lerman

*Scientific Reports*

### **Pyleoclim: Paleoclimate Timeseries Analysis and Visualization with Python**

D. Khider, J. Geay, F. Zhu, A. James, J. Landers, V. Ratnakar, and Y. Gil

*Paleoceanography and Paleoclimatology*

### **Quantifying How Hateful Communities Radicalize Online Users**

M. Schmitz, K. Burghardt, and G. Muric

*ASONAM 2022*

### **Reflect, Not Reflex: Inference-Based Common Ground Improves Dialogue Response Quality**

P. Zhou, H. Cho, P. Jandaghi, D. Lee, B. Lin, J. Pujara, and X. Ren

*Conference on Empirical Methods in Natural Language Processing*

### **Research Gaps and Opportunities in Precision Nutrition: An NIH Workshop Report**

B. Lee, J. Ordovás, E. Parks, C. Anderson, A. Barabási, S. Clinton, K. de la Haye, V. Duffy, P. Franks, E. Ginexi, K. Hammond, E. Hanlon, M. Hittle, E. Ho, A. Horn et al.

*American Journal of Clinical Nutrition*

### **Residual Learning to Integrate Neural Network and Physics-Based Models for Improved Production Prediction in Unconventional Reservoirs**

J. Cornelio, S. Mod Razak, Y. Cho, H. Liu, R. Vaidya, and B. Jafarpour

*SPE Journal*

### **Road Network Evolution in the Urban and Rural United States since 1900**

K. Burghardt, J. Uhl, K. Lerman, and S. Leyk

*Computers, Environment and Urban Systems*

### **Scale-Free Degree Distributions, Homophily and the Glass Ceiling Effect in Directed Networks**

B. Nettasinghe, N. Alipourfard, S. Iota, V. Krishnamurthy, and K. Lerman

*Journal of Complex Networks*

### **Secure Domain Adaptation with Multiple Sources**

S. Stan and M. Rostami

*Transactions on Machine Learning Research*

### **The Social Emotional Web**

K. Lerman

*Proceedings of the IEEE International Conference on Collaboration and Internet Computing (IEEE CIC)*

### **Think Before You Speak: Explicitly Generating Implicit Commonsense Knowledge for Response Generation**

P. Zhou, K. Gopalakrishnan, B. Hedayatnia, S. Kim, J. Pujara, X. Ren, Y. Liu, and D. Hakkani-Tur

*Association for Computational Linguistics*

### **TILES-2019: A Longitudinal Physiologic and Behavioral Data Set of Medical Residents in an Intensive Care Unit**

J. Yau, B. Girault, T. Feng, K. Mundnich, A. Nadarajan,

B.M. Booth, E. Ferrara, K. Lerman, E. Hsieh, and S. Narayanan

*Scientific Data*

### **Towards Capturing Scientific Reasoning to Automate Data Analysis**

Y. Gil, D. Khider, M. Osorio, V. Ratnakar, H. Vargas, D. Garijo, and S. Pierce

*Proceedings of the 44th Annual Conference of the Cognitive Science Society (CogSci)*

### **Transfer Learning via Representation Learning**

M. Rostami, H. He, M. Chen, and D. Roth

*Federated and Transfer Learning*

### **Transfer Learning with Multiple Aggregated Source Models in Unconventional Reservoirs**

J. Cornelio, S. Mohd Razak, Y. Cho, H.H. Liu, R. Vaidya, and B. Jafarpour

*Unconventional Resources Technology Conference*

### **Transfer Learning with Recurrent Neural Networks for Long-Term Production Forecasting in Unconventional Reservoirs**

S. Mohd Razak, J. Cornelio, Y. Cho, H.H. Liu, R. Vaidya, and B. Jafarpour

*SPE Journal*

### **Understanding Narratives through Dimensions of Analogy**

T. Nagarajah, F. Ilievski, and J. Pujara

*IJCAI Workshop on Qualitative Reasoning*

### **Unequal Impact and Spatial Aggregation Distort Covid-19 Growth Rates**

K. Burghardt, S. Guo, and K. Lerman

*Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*

**Visual Sudoku Puzzle Classification: A Suite of Collective Neuro-Symbolic Tasks**

E. Augustine, C. Pryor, C. Dickens, J. Pujara, W. Wang, and L. Getoor

*Workshop on Neural Symbolic Learning and Reasoning*

**Where Does Bias in Common Sense Knowledge Models Come From?**

S. Melotte, F. Ilievski, L. Zhang, A. Malte, N. Mutha, F. Morstatter, and N. Mehrabi

*IEEE Internet Computing*

**Zero-Shot Image Classification using Coupled Dictionary Embedding**

M. Rostami, S. Kolouri, Z. Murez, Y. Owechko, E. Eaton, and K. Kim

*Machine Learning with Applications*

**PaleoRec: A Sequential Recommender System for the Annotation of Paleoclimate Datasets**

S. Manety, D. Khider, C. Heiser, N. McKay, J. Emile-Geay, and C. Routson

*Environmental Data Science*

**PipeEdge: Pipeline Parallelism for Large-Scale Model Inference on Heterogeneous Edge Devices**

Y. Hu, C. Imes, X. Zhao, S. Kundu, P. Beerel, S. Crago, and J.P. Walters

**GCN-WP — Semi-Supervised Graph Convolutional Networks for Win Prediction in Esports**

A.J. Bisberg and E. Ferrara

*2022 IEEE Conference on Games (CoG)*

**The Gift That Keeps on Giving: Generosity is Contagious in Multiplayer Online Games**

A.J. Bisberg, J. Jiang, Y. Zeng, E. Chen, and E. Ferrara

*CSCW 2022*

**Comparative Analysis of Social Bots and Humans During the COVID-19 pandemic**

H.C.H. Chang and E. Ferrara

*Journal of Computational Social Science*

**Construction of Large-Scale Misinformation Labeled Datasets from Social Media Discourse using Label Refinement**

K. Sharma, E. Ferrara, and Y. Liu

*WWW '22: Proceedings of the ACM Web Conference 2022*

**Botometer 101: Social Bot Practicum for Computational Social Scientists**

K.C. Yang, E. Ferrara, and F. Menczer

*Journal of Computational Social Science*

**Twitter Spam and False Accounts Prevalence, Detection and Characterization: A Survey**

E. Ferrara

*First Monday*

**#JusticeForGeorgeFloyd: How Instagram Facilitated the 2020 Black Lives Matter Protests**

H.C.H. Chang, A. Richardson, and E. Ferrara

*PLOS ONE*

**Characterizing Online Engagement with Disinformation and Conspiracies in the 2020 US Presidential Election**

K. Sharma, E. Ferrara, and Y. Liu

*ICWSM 2022 - 16th International AAAI Conference on Web and Social Media*

**Microelectronics and Electronics Systems**

**A Physics Based MTJ Compact Model for State-of-the-Art and Emerging STT-MRAM Failure Analysis and Yield Enhancement**

N. Gaul, A. Jaiswal, H. Yoon, T. Lee, K. Yamane, J. Versaggi, R. Carter, and C. Paul

*2022 IEEE International Memory Workshop (IMW)*

**A Processing-in-Pixel-in-Memory Paradigm for Resource-Constrained TinyML Applications**

G. Datta, S. Kundu, Z. Yin, R. Lakkireddy, J. Mathai, A. Jacob, P. Beerel, and A. Jaiswal

*Scientific Reports*

**ACE-SNN: Algorithm-Hardware Co-Design of Energy-Efficient & Low-Latency Deep Spiking Neural Networks for 3D Image Recognition**

G. Datta, S. Kundu, A. Jaiswal, and P. Beerel

*Frontiers in Neuroscience*

**Coupled Microdisk Resonators for Efficient Electro-Optic Quantum Frequency Conversion**

R. Kudalippallyalil, S. Chandran, A. Jaiswal, and A. Jacob

*2022 Conference on Lasers and Electro-Optics (CLEO)*

**CryoCiM: Cryogenic Compute-in-Memory Based on the Quantum Anomalous Hall Effect**

S. Alam, M. Islam, M. Hossain, A. Jaiswal, and A. Aziz

*Applied Physics Letters*

**Heterogeneously Integrated Quantum Chip Interposer Packaging**

R. Kudalippallyalil, S. Chandran, A. Jaiswal, K. Wang, and A. Jacob

*2022 IEEE 72nd Electronic Components and Technology Conference (ECTC)*

**Modeling the Energy Efficiency of GEMM using Optical Random Access Memory**

B. Zhang, A. Jaiswal, C. Mathew, R. Lakkireddy, A. Jacob, S. Wijeratne, and V. Prasanna

*2022 IEEE High Performance Extreme Computing Conference (HPEC)*

**P 2 M-DeTrack: Processing-in-Pixel-in-Memory for Energy-efficient and Real-Time Multi-Object Detection and Tracking**

G. Datta, S. Kundu, Z. Yin, J. Mathai, Z. Liu, Z. Wang, M. Tian, S. Lu, R. Lakkireddy, A. Schmidt et al.  
2022 IFIP/IEEE 30th International Conference on Very Large Scale Integration (VLSI-SoC)

**Performance Modeling Sparse MTTKRP Using Optical Static Random Access Memory on FPGA**

S. Wijeratne, A. Jaiswal, A. Jacob, B.Zhang, and V. Prasanna  
2022 IEEE High Performance Extreme Computing Conference (HPEC)

**Variation-Aware Design Space Exploration of Mott Memristor-Based Neuristors**

S. Alam, M. Islam, A. Jaiswal, N. Cady, G. Rose, and A. Aziz  
2022 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)

**Natural Language Processing**

**An Improved Baseline for Sentence-Level Relation Extraction**

W. Zhou and M. Chen  
*Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (AACL)*

**Answer Consolidation: Formulation and Benchmarking**

W. Zhou, Q. Ning, H. Elfardy, K. Small, and M. Chen  
*Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*

**Augmenting Training Data for Massive Semantic Matching Models in Low-Traffic E-commerce Stores**

A. Joshi, S. Vishwanath, C. Teo, V. Petricek, V. Vishwanathan, R. Bhagat, and J. May  
*Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics Human Language Technologies Industry Track*

**Bending the Future: Autoregressive Modeling of Temporal Knowledge Graphs in Curvature-Variable Hyperbolic Spaces**

J. Sohn, M. Ma, and M. Chen  
*4th Conference on Automated Knowledge Base Construction (AKBC)*

**Building an Event Extractor with Only a Few Examples**

P Yu, Z. Zhang, C. Voss, H. Ji, and J. May  
*Proceedings of the Third Workshop on Deep Learning for Low-Resource Natural Language Processing*

**Contextualized Scene Imagination for Generative Commonsense Reasoning**

P. Wang, J. Zamora, J. Liu, F. Ilievski, M. Chen, and X. Ren  
*International Conference on Learning Representations*

**Dangling-Aware Entity Alignment with Mixed High-Order Proximities**

J. Liu, Z. Sun, B. Hooi, Y. Wang, D. Liu, B. Yang, X. Xiao, and M. Chen  
*Findings of the Association for Computational Linguistics: NAACL 2022*

**Distant Supervision for Preconditioned Commonsense Inference**

E. Qasemi, P. Khanna, Q. Ning, and M. Chen  
*SoCal NLP Symposium*

**Does Your Model Classify Entities Reasonably? Diagnosing and Mitigating Spurious Correlations in Entity Typing**

N. Xu, F. Wang, B. Li, M. Dong, and M. Chen  
*Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*

**GRAPHCACHE: Message Passing as Caching for Sentence-Level Relation Extraction**

Y. Wang, M. Chen, W. Zhou, Y. Cai, Y. Liang, and B. Hooi  
*Findings of the Association for Computational Linguistics: NAACL 2022*

**Know Thy Strengths: Comprehensive Dialogue State Tracking Diagnostics**

H. Cho, C. Sankar, C. Lin, K. Sadagopan, S. Shayandeh, A. Celikyilmaz, J. May, and A. Beirami  
*Findings of the Association for Computational Linguistics: EMNLP 2022*

**Language Model Priming for Cross-Lingual Event Extraction**

S. Fincke, S. Agarwal, S. Miller, and E. Boschee  
*Proceedings of the AAAI Conference on Artificial Intelligence*

**Machine Translation Robustness to Natural Asemantic Variation**

J. Bremerman, X. Ren, and J. May  
*Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*

**New Frontiers of Information Extraction**

M. Chen, L. Huang, M. Li, B. Zhou, H. Ji, and D. Roth  
*Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Tutorials*

**Opponent Modeling in Negotiation Dialogues by Related Data Adaptation**

K. Chawla, G. Lucas, J. May, and J. Gratch  
*Findings of the Association for Computational Linguistics NAACL2022*

**PaCo: Preconditions Attributed to Commonsense Knowledge**

E. Qasemi, F. Ilievski, M. Chen, and P. Szekely  
*Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*

**PlnKS: Preconditioned Commonsense Inference with Minimal Supervision**

E. Qasemi, P. Khanna, Q. Ning, and M. Chen  
*Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing*

**Prix-LM: Pretraining for Multilingual Knowledge Base Construction**

W. Zhou, F. Liu, I. Vulić, N. Collier, and M. Chen  
*Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*

**Robust (Controlled) Table-to-Text Generation with Structure-Aware Equivariance Learning**

F. Wang, Z. Xu, P. Szekely, and M. Chen  
*Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*

**Saliency Allocation as Guidance for Abstractive Summarization**

F. Wang, K. Song, H. Zhang, L. Jin, S. Cho, W. Yao, X. Wang, M. Chen, and D. Yu  
*Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*

**Segmenting Numerical Substitution Ciphers**

N. Aldarrab and J. May  
*Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*

**Sharpness-Aware Minimization with Dynamic Reweighting**

W. Zhou, F. Liu, H. Zhang, and M. Chen  
*Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*

**Should We Rely on Entity Mentions for Relation Extraction? Debiasing Relation Extraction with Counterfactual Analysis**

Y. Wang, M. Chen, W. Zhou, Y. Cai, Y. Liang, D. Liu, B. Yang, J. Liu, and B. Hooi  
*Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*

**SOSum: a Dataset of Stack Overflow Post Summaries**

B. Kou, Y. Di, M. Chen, and T. Zhang  
*Proceedings of the 19th International Conference on Mining Software Repositories*

**SpaBERT: A Pretrained Language Model from Geographic Data for Geo-Entity Representation**

Z. Li, J. Kim, Y. Chiang, and M. Chen  
*Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*

**Summarization as Indirect Supervision for Relation Extraction**

K. Lu, I. Hsu, Zhou, Wenxuan and M. Ma, M. Chen et al.  
*Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*

**Transfer-Based Taxonomy Induction over Concept Labels**

M. Kejriwal, K. Shen, C. Ni, and N. Torzec  
*Engineering Applications of Artificial Intelligence*

**Ultra-Fine Entity Typing with Indirect Supervision from Natural Language Inference**

B. Li, W. Yin, and M. Chen  
*Transactions of the Association for Computational Linguistics*

**Unified Semantic Typing with Meaningful Label Inference**

J. Huang, B. Li, J. Xu, and M. Chen  
*Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*

**DEGREE: A Data-Efficient Generation-Based Event Extraction Model**

I. Hsu, K. Huang, E. Boschee, S. Miller, P. Natarajan, K Chang, and N. Peng  
*Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*

**NewsEdits: A News Article Revision Dataset and a Novel Document-Level Reasoning Challenge**

A. Spangher, X. Ren, J. May, and N. Peng  
*Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*

**Networking**

**A FastMap-Based Algorithm for Block Modeling**

A. Li, P. Stuckey, S. Koenig, and S. Thittamaranahalli  
*Proceedings of the Nineteenth International Conference on the Integration of Constraint Programming, Artificial Intelligence and Operations Research (CPAIOR-2022)*

**A\*pex: Efficient Approximate Multi-Objective Search on Graphs**

H. Zhang, O. Salzman, S. Thittamaranahalli, A. Felner, C. Hernandez, and S. Koenig  
*Proceedings of the Thirty-Second International Conference on Automated Planning and Scheduling (ICAPS-2022)*

**AI-Driven Quantification of Ground Glass Opacities in Lungs of COVID-19 Patients Using 3D Computed Tomography Imaging**

M. Saha, S. Amin, A. Sharma, S. Thittamaranahalli, and R. Kalia  
*PLOS ONE Open Access Journal 17(3)*

### AmLight Real Time In-Band Telemetry Within the NoviFlow Switches

J. Bezerra, J. Ibarra, A. Torres, I. Brito, V. Chergarova, H. Morgan, M. LeClerc, and B. Ker

*Future of Information and Communication Conference*

### Anycast Agility: Network Playbooks to Fight DDoS

ASM Rizvi, L. Bertholdo, J. Ceron, and J. Heidemann

*Proceedings of the 31st USENIX Security Symposium*

### Anytime Approximate Bi-Objective Search

H. Zhang, O. Salzman, S. Thittamaranahalli, A. Felner, C. Hernandez, and S. Koenig

*Proceedings of the Fifteenth International Symposium on Combinatorial Search (SOCS-2022)*

### App2Net: Moving Application Functions to Network & a Case Study on Low-Latency Feedback

G. Sankaran, J. Chung, and R. Kettimuthu

*2022 IEEE/ACM International Workshop on Innovating the Network for Data-Intensive Science (INDIS)*

### Chhoyhopper: A Moving Target Defense with IPv6

A. Rizvi and J. Heidemann

*Proceedings of the IEEE Workshop on Measurements*

### Considerations for Large Authoritative DNS Server Operators (RFC)

G. Moura, W. Hardaker, J. Heidemann, and M. Davids

*Internet Engineering Task Force, Internet Society*

### Content Delivery Networks (CDNs) Overview and Illicit Material Considerations

W. Hardaker

[publicexchange.usc.edu](http://publicexchange.usc.edu)

### Cybersecurity as Illuminator for the Future of Computing Research

J. Wroclawski and T. Benzel

*Communications of the ACM*

### Differences in Monitoring the DNS Root Over IPv4 and IPv6

T. Saluja, J. Heidemann, and Y. Pradkin

*Proceedings of the National Symposium for NSFREU Research in Data Science*

### Explaining Deep Learning Models for Per-Packet Encrypted Network Traffic Classification

L. Garcia, G. Bartlett, S. Ravi, H. Ibrahim, W. Hardaker, and E. Kline

*2022 IEEE International Symposium on Measurements & Networking (M&N)*

### FastMapSVM: Classifying Seismograms Using the FastMap Algorithm and Support-Vector Machines

M. White, K. Sharma, A. Li, S. Thittamaranahalli, and N. Nakata

*Seismological Research Letters 93(2B)*

### FastPivot: An Algorithm for Inverse Problems

Y. Guan, A. Li, S. Koenig, S. Haas, and S. Thittamaranahalli

*Proceedings of the Eighteenth IEEE International Conference on Automation Science and Engineering (CASE- 2022)*

### Improving Fidelity in Video Streaming Experimentation on Testbeds with a CDN

C. Ardi, A. Hussain, M. Collins, and S. Schwab

*Proceedings of the 2nd International Workshop on Design, Deployment, and Evaluation of Network-Assisted Video Streaming*

### Internet Outage Detection using Passive Analysis (poster abstract)

A. Enayet and J. Heidemann

*Proceedings of the ACM Internet Measurement Conference*

### Multi-Agent Path Finding with Mutex Propagation

H. Zhang, J. Li, P. Surynek, S. Thittamaranahalli, and S. Koenig

*Artificial Intelligence Journal (AIJ-2022)*

### Mutex Propagation in Multi-Agent Path Finding for Large Agents

H. Zhang, Y. Li, J. Li, S. Thittamaranahalli, and S. Koenig

*Proceedings of the Fifteenth International Symposium on Combinatorial Search (SOCS-2022)*

### Old but Gold: Prospecting TCP to Engineer and Live Monitor DNS Anycast

G. Moura, J. Heidemann, W. Hardaker, P. Charnsethikul,

J. Bulten, J. Ceron, and C. Hesselman

*Proceedings of the Passive and Active Measurement Workshop*

### Securing 5G Slices using Homomorphic Encryption

E. Kline, S. Ravi, D. Cousins, and S. Rv

*2022 IEEE Wireless Communications and Networking Conference (WCNC)*

### Towards Production Deployment of a SDX Control Framework

M. Cevik, M. Stealey, C. Wang, J. Bezerra, J. Ibarra,

V. Chergarova, H. Morgan, and Y. Xin

*International Conference on Computer Communications and Networks (ICCCN)*

### Trajectory Optimization for Safe Navigation in Maritime Traffic Using Historical Data

B. Chaithanya, A. Singh, A. Sinha, A. Kumar,

and S. Thittamaranahalli

*Proceedings of the Twenty-Eighth International Conference on Principles and Practice of Constraint Programming (CP-2022)*



## Quantum Information Sensing, Communication and Processing

### A Long Wave Infrared Communication Transmitter using Tunable Emissivity in Graphene

H. Weinstein, J. Habif, Z. Cai, and S. Cronin  
*Frontiers in Optics*

### Demonstration of Quantum-Limited Discrimination of Multicopy Pure Versus Mixed States

A. Jagannathan, M. Grace, O. Brasher, J. Shapiro, S. Guha, and J. Habif  
*Physical Review A*

### Electrically Modulated Thermal Emissivity in Multi-Layer Graphene Devices

J. Habif, S. Cronin, A. Jagannathan, Z. Cai, and H. Weinstein  
*Bulletin of the American Physical Society*

### Experimental Demonstration of an Optical Half-Adder of Two 4-PSK, 10-Gbit/s Channels using Nonlinear Wave Mixing

H. Song, K. Zou, N. Karapetyan, A. Minoofar, H. Zhou, X. Su, A. Almainan, J. Habif et al.  
*2022 IEEE Photonics Conference (IPC)*

### Experimental Demonstration of Remotely Controlled Tunable Optical Correlators of 10--50 Gbaud QPSK Channels using Linear and Nonlinear Components and Laser-Delivered Powers

F. Alishahi, A. Minoofar, A. Fallahpour, K. Zou, N. Karapetyan, H. Zhou, J. Habif et al.  
*Optics Communications*

### Gate-Tunable Modulation of the Optical Properties of Multilayer Graphene by the Reversible Intercalation of Ionic Liquid Anions

Z. Cai, I. Aravind, H. Weinstein, R. Li, J. Wu, H. Wang, J. Habif, and S. Cronin  
*Journal of Applied Physics*

### Harvesting Planck Radiation for Free-Space Optical Communications in the Long-Wave Infrared Band

H. Weinstein, Z. Cai, S. Cronin, and J. Habif  
*Optics Letters*

### Integrated All-Optical Controlled Amplitude Modulator Using Laser Delivered Signals

S. Idres, A. Fallahpour, A. Willner, J. Habif, and H. Hashemi  
*CLEO: Applications and Technology*

### Quantum Annealing with Special Drivers for Circuit Fault Diagnostics

H. Leipold and F. Spedalieri  
*Scientific Reports*

### Quantum Bounds for Discriminating Mixed States Generated by Weak Measurements and Thermal Noise

P. Wysocki, J. Habif, and T. McAskill  
*Bulletin of the American Physical Society*

### Remotely Biasing, Controlling, and Monitoring a Network Routing Node Based on Optically Provided Signals

A. Minoofar, F. Alishahi, A. Fallahpour, K. Zou, N. Karapetyan, J. Habif, and A. Willner  
*Metro and Data Center Optical Networks and Short-Reach Links V*

## Scientific Computing

### Accelerating Scientific Workflows on HPC Platforms with In Situ Processing

T. Do, L. Pottier, O. Yildiz, K. Vahi, P. Krawczuk, T. Peterka, and E. Deelman  
*2022 IEEE/ACM 22nd International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*

### Automating Edge-to-Cloud Workflows for Science: Traversing the Edge-to-Cloud Continuum with Pegasus

R. Tanaka, G. Papadimitriou, S. Viswanath, C. Wang, E. Lyons et al.  
*2022 22nd IEEE International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*

### Broadening Student Participation in Cyberinfrastructure Research and Development

A. Murillo, E. Deelman, J. Nabrzyski, and L. Pottier  
*Practice and Experience in Advanced Research Computing*

### Building the Research Innovation Workforce: Challenges and Recommendations from a Virtual Workshop to Advance the Research Computing Community\*

T. Hacker, P. Smith, D. Brunson, L. Arafune, T. Cheatham, and E. Deelman  
*Practice and Experience in Advanced Research Computing*

### Co-Scheduling Ensembles of In Situ Workflows

T. Do, L. Pottier, R. Ferreira da Silva, F. Suter, S. Caño-Lores, M. Taufer, and E. Deelman  
*2022 IEEE/ACM Workshop on Workflows in Support of Large-Scale Science (WORKS)*

### Data Integrity Error Localization in Networked Systems with Missing Data

Y. Xin, S. Fu, A. Mandal, R. Tanaka, M. Rynge, K. Vahi, and E. Deelman  
*ICC 2022 - IEEE International Conference on Communications*

### Managing Database-Application Co-Evolution in a Scientific Data Ecosystem

R. Schuler and C. Kesselman  
*2022 IEEE 18th International Conference on E-Science (e-Science)*

### Novel Proposals for FAIR, Automated, Recommendable, and Robust Workflows

I. Abhinit, E. Adams, K. Alam, B. Chase et al.  
*2022 IEEE/ACM Workshop on Workflows in Support of Large-Scale Science (WORKS)*

### **Performance Assessment of Ensembles of In Situ Workflows under Resource Constraints**

T. Do, L. Pottier, R. Ferreira da Silva et al.  
*Concurrency and Computation: Practice and Experience*

### **The Multi-Tier Assistance, Training, and Computational Help (MATCH) Project, a Track 2 NSF ACCESS Initiative**

S.L. Knuth, J. Ma, J.C. Adams, A. Chalker, E. Deelman et al.  
*The Journal of Computational Science Education*

### **Regional Synapse Gain and Loss Accompany Memory Formation in Larval Zebrafish**

W. Dempsey, Z. Du, A. Nadtochiy, C. Smith, K. Czajkowski, A. Andreev, D. Robson, J. Li, S. Applebaum, T. Truong, C. Kesselman, S. Fraser, and D. Arnold  
*Proceedings of the National Academy of Sciences*

### **VisDict: A Visual Dictionary in a Science Gateway**

S. Gesing, E. Deelman, R. Ferreira da Silva et al.  
*Proceedings, Gateways 2022*

### **WfCommons: A Framework for Enabling Scientific Workflow Research and Development**

T. Coleman, H. Casanova, and L. Pottier et al.  
*Future Generation Computer Systems*

### **Workflow Anomaly Detection with Graph Neural Networks**

H. Jin, K. Raghavan, G. Papadimitriou et al.  
*2022 IEEE/ACM Workshop on Workflows in Support of Large-Scale Science (WORKS)*

### **Making Common Fund Data More Findable: Catalyzing a Data Ecosystem**

A. Charbonneau, A. Brady, K. Czajkowski et al.  
*GigaScience*

## **Vision, Image, Speech, and Text Analytics**

### **Attack-Agnostic Adversarial Detection**

J. Cheng, M. Hussein, J. Billa, and W. AbdAlmageed  
*NeurIPS 2022 Workshop on Trustworthy and Socially Responsible Machine Learning (TSRML)*

### **August: Automating Usage-Based Test Generation from Videos of App Executions**

Y. Zhao, S. Talebipour, K. Baral, H. Park, L. Yee et al.  
*Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*

### **Intelligent Traffic Monitoring with Hybrid AI**

E. Qasemi and A. Oltramari  
*IJCAI Workshop on Artificial Intelligence for Autonomous Driving (AI4AD)*

### **Understanding Multimodal Procedural Knowledge by Sequencing Multimodal Instructional Manuals**

T. Wu, A. Spangher, P. Alipoormolabashi, M. Freedman, R. Weischedel, and N. Peng  
*Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*



# CLOUDWALKERS

## ISI AND THE INVENTORS OF THE FUTURE

USC VITERBI PRESENTS "CLOUDWALKERS ISI AND THE INVENTORS OF THE FUTURE" PRODUCED BY ADAM TYREE DAVID BEIER  
CO-PRODUCER ROY TIGHE EXECUTIVE PRODUCED BY YANNIS YORTSOS CRAIG KNOBLOCK CINEMATOGRAPHY BY JONATHAN POPE  
EDITED BY ROMMEL VILLA BARRIGA PRODUCTION DESIGN BY ANCA COCAN SOUND BY ERIC HAN  
WRITTEN, DIRECTED, AND PRODUCED DANIEL DRUHOVA

USC Viterbi  
School of Engineering





4676 Admiralty Way #1001 | Marina del Rey, CA 90292 | 310.822.1511 | [www.isi.edu](http://www.isi.edu) | [media@isi.edu](mailto:media@isi.edu)

---

**USC Viterbi**  
School of Engineering