

MARCH
16-17, 2018

SAN DIEGO

GRAD Cohort URMD

Underrepresented Minorities & Persons with Disabilities

2018



CRA

Computing Research
Association



Dear Grad Cohort Participant,

We welcome you to the 2018 CRA Grad Cohort Workshop for Underrepresented Minorities + Persons with Disabilities (URMD)! The next few days are filled with sessions where 25 senior computing researchers and professionals will be sharing their strategies and experiences to help increase your graduate school and career success. There will also be plenty of opportunities to meet and network with these successful researchers as well as with graduate students from other universities. We hope that you will take the utmost advantage of this unique experience by actively participating in discussions, developing peer networks, and building mentoring relationships.

Since this is the inaugural URMD Grad Cohort Workshop, we are especially interested in hearing your feedback about the program and the experience. Please take time to complete the evaluation form provided after the workshop. We want to learn what you liked and did not like, as well as any suggestions you might have for improving the event in subsequent years.

The 2018 CRA-URMD Workshop is made possible through generous contributions by the Computing Research Association, National Science Foundation, AccessComputing, Whova, Google and Association for Computing Machinery. Please join us in thanking them for their kind support.

We hope that you take home many new insights and connections from this workshop to help you succeed in your graduate school and research career. Be ready to be inspired, learn, and meet many interesting researchers in the field.

Sincerely,

Lori Clarke, Ayanna Howard, Kunle Olukotun, Mary Lou Soffa
CO-CHAIRS, CRA-URMD GRADUATE COHORT WORKSHOP




SAN DIEGO, CA
MARCH 16-17, 2018
Omni San Diego

2018 GradCohort

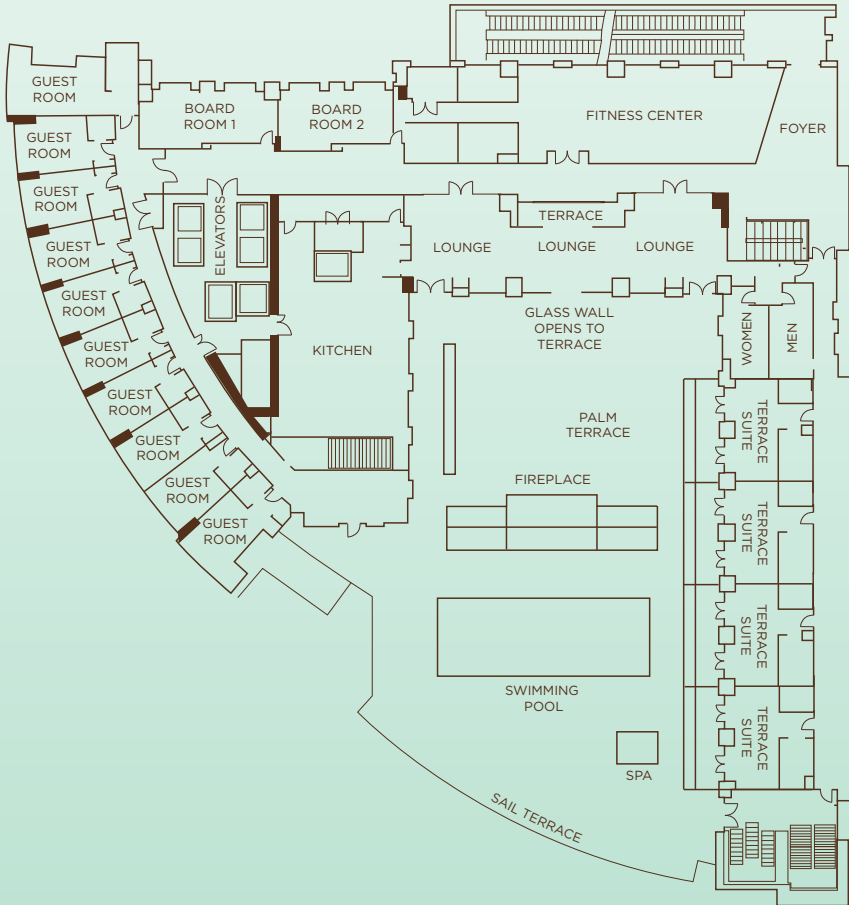


Floor Plan

San Diego Gallery Meeting Space



San Diego 6th Floor





Program



Thursday, March 15, 2018

1:00pm - 6:30pm **Early Registration - Hotel Lobby**

6:00pm - 8:00pm **Welcome Reception - Palm Terrace**

Friday, March 16, 2018

7:00am - 8:45am **Registration - Art Foyer**

7:30am - 8:30am **Breakfast - Gallery 1**

8:30am - 8:45am **Welcome - Gallery 2**

Speaker(s): Lori Clarke, Ayanna Howard, Kunle Olukotun & Mary Lou Soffa

	Early Years <i>Gallery 2</i>	Middle Years <i>Gallery 3A</i>	Later Years <i>Gallery 3B</i>
8:50am - 9:50am	Networking	Finding a Research Topic & Interdisciplinary Research	Industry vs. Academic Research Positions
Speaker(s):	Jeanine Cook & James Mickens	Melanie Moses & Armando Solar-Lezama	Dilma Da Silva
9:55am - 10:55am	Master's vs. Ph.D.	Presentation and Other Verbal Communication Skills	Preparing Your Thesis Proposal and Becoming a Ph.D. Candidate
Speaker(s):	Russ Joseph & CJ Taylor	Shiri Azenkot & Ramon Caceres	Dorian Arnold & Shaun Kane
10:55am - 11:15am Break - Art Foyer			
11:15am - 12:15pm	M.S. Career Opportunities and Job Search	Finding the Help You Need and Deserve	Ph.D. Academic Career Paths and Job Search
Speaker(s):	Jeanine Cook & Jaime Moreno	Jeff Forbes & Lydia Tapia	Chad Jenkins & Armando Solar-Lezama
12:15pm - 1:30pm ACM Sponsored Lunch (Research themed lunch) - Sail Terrace			
1:30pm - 2:30pm	Finding an Advisor and Developing an Effective Working Relationship with Them (and what to do if it doesn't work)	Balancing Graduate School and Personal Life	Ph.D. Non-Academic Career Paths and Job Search
Speaker(s):	Jason Mars & Melanie Moses	Jaime Moreno & CJ Taylor	Alan Bivens & Daniel Lustig
2:30pm - 3:00pm Break			

3:00pm - 4:00pm **Finding Your Way** (Overcoming cultural barriers) - *Gallery 2*

Speaker(s): Ayanna Howard & Daniela Marghitu

4:00pm - 4:15pm **Break** - *Art Foyer*

4:15pm - 5:00pm

Poster Session

BOF - Overcoming Insufficient Academic Preparation: Perceived and Real
Dilma Da Silva

BOF - Empowerment of People with Disabilities
Daniela Marghitu

5:00pm - 5:45pm **Poster Session**

6:30pm - 10:30pm **Reception** - *Palm Terrace*

Saturday, March 17, 2018

7:30am - 8:30am **Breakfast** - *Gallery 1*

8:30am - 9:40am **Strategies for Human-Human Interaction** - *Gallery 2*

Speaker(s): Shiri Azenkot, Dorian Arnold, and Daniela Marghitu

9:40am - 10:10am **Break** - *Art Foyer*

10:10am - 11:10am
Speaker(s):

Early Years
Gallery 2

Middle Years
Gallery 3A

Later Years
Gallery 3B

Financing Your Graduate Education

Building Your Professional Persona

Entrepreneurship Opportunities & Skills

Jeff Forbes

Chad Jenkins &
Daniel Lustig

Jason Mars &
Kunle Olukotun

11:20am - 12:20pm
Speaker(s):

Summer Internships
Alan Bivens &
Russ Joseph

Building Self-Confidence
Shaun Kane &
Lydia Tapia

Publishing Your Research
Ramon Caceres &
Dilma Da Silva

12:30pm - 12:45pm **Wrap-Up & Final Remarks** - *Gallery 2*

Speaker(s): Lori Clarke, Ayanna Howard, Kunle Olukotun & Mary Lou Soffa

12:45pm - 2:00pm **Lunch** - *Sail Terrace*

2:00pm - 4:00pm
Speaker(s):

Individual Resume/CV Advising -
Gallery 3A

Individual Academic/Career Advising -
Gallery 3B



URMD Grad Cohort Session Descriptions

Early Years

► Networking

This session will address the skills that are needed for networking - a very important component of your professional life. The topics include strategies for finding a community, meeting people in the field, and promoting your research and yourself. This session will also address what you should prepare for and what to do when you attend conferences, workshops, or any technical meeting, including meeting researchers visiting your department.

► Master's vs. Ph.D.

This session will address the main differences in career options and graduate studies between a Master's and a Ph.D. degree and how to identify financial support for your academic studies. Topics include realistic goals and expectations for each degree and strategies for deciding which degree to pursue.

► M.S. Career Opportunities and Job Search

This session will be a discussion of the various career opportunities for Master's graduates and how to find a position that is best for you. Discussion will include possible career paths and the role of mentors in industry. This session will also discuss the process of finding an industry position for M.S. graduates, how to prepare for the interview, what questions to ask and what you should expect during the interview, and what to do after the interview. Tips will be given on writing a cover letter and resume that attracts attention.

► Financing Your Graduate Education

This session will focus on how to identify financial support for your academic studies beyond the initial teaching/research assistantship position, including positioning yourself for research assistantships, tracking down fellowship opportunities, and how to put together a successful application. This session will also address how to survive temporary funding shortfalls/lapses.

► Finding an Advisor and Developing an Effective Working Relationship

Obtaining a Ph.D. is like undertaking an apprenticeship. Thus, it is very important to have a good advisor and a good relationship with that advisor. This session will focus on the importance of carefully choosing an advisor, how to get the most out of your

interactions with your advisor, and the responsibilities of both student and advisor in making the graduate research experience successful. It will discuss how to work through problems with your advisor, when to consider changing advisors or institutions, and strategies for undertaking such a change. It will also discuss the benefits of having mentors, in addition to your advisor.

► Summer Internships

This session will focus on how to go about finding an industry or government lab internship for the summer, the advantages and disadvantages of internships, and how to get the most out of your summer experience.

Mid Years

► Finding a Research Topic & Interdisciplinary Research

This session will focus on strategies for actively identifying a viable research topic for your Ph.D. dissertation. It will discuss how to choose between several topics, and how to balance scientific curiosity with professional demands. It will address how your career plans (e.g., teaching, academic research, or industrial research or development) may influence your choice. It will also describe the various aspects of a research topic (e.g., theoretical foundation, experimental framework, analytic and experimental results), and it will discuss how to set sensible goals and milestones so that you can successfully finish in a reasonable time frame.

► Presentation and Other Verbal Communication Skills

This session will focus on building your oral communication skills. Topics include strategies for high quality oral presentations of papers, posters, and panels.

► Finding the Help You Need and Deserve

► **Balancing Graduate School and Personal Life**

This session will address strategies for maintaining balance and a positive perspective in your life during graduate school and throughout your career. Topics will include achieving personal satisfaction in your career and family life, balancing your TA duties, your course work, and your research program, decision-making strategies when career and family priorities seem to collide, and re-evaluating your personal and career goals and initiating changes.

► **Building Your Professional Persona**

This session will address the dos and don'ts of building a professional image. Topics will include web presence (e.g., personal pages and social media), dissemination of technical contributions, and professional ethics.

► **Building Self-Confidence**

This session will address the confidence crisis that 2nd (and 3rd) year graduate students often face: how to recover from not doing as well in a course as you expected, from not passing the Ph.D. candidacy exams on your first try, from getting a paper rejected, from the frustration of not knowing what your specific research project will be, and from the feeling that you don't know as much as your fellow graduate students. The discussion will focus on issues and rewards of continuing towards your goals and milestones in grad school and mustering the courage to persist.

Final Years

► **Industry vs. Academic Research Positions**

This session will focus on the difference in research positions and opportunities in industry and academia. It will describe the expectations and challenges in both and considerations in choosing between the two, including opportunities for moving from one to the other.

► **Preparing Your Thesis Proposal and Becoming a Ph.D. Candidate**

This session will cover preparing a thesis proposal. A strong thesis proposal demonstrates an understanding of related literature and lays out a plan for completing your thesis. This session will discuss the goals of a thesis proposal, an outline of what to include, strategies to choose your committee, and how to prepare for your proposal defense.

► **Ph.D. Academic Career Paths and Job Search**

This session will focus on the different career paths in academia. Topics include the roles of research, teaching, and service, and

how they differ in different academic institutions as well as depending on one's position (e.g., tenure track faculty, lecturer, administrator). This session will also discuss the challenges and rewards of research, teaching, service, and mentoring undergraduate students and the skills for success in each. It will also discuss changing career paths.

► **Ph.D. Non-Academic Career Paths and Job Search**

This session will focus on the different career paths for new Ph.D.'s in industry and government laboratories. Topics will include the challenges and rewards of a non-academic career, skills and experiences needed for success, opportunities for advancement and taking the initiative, collaborating with researchers in academia or other organizations, and positioning yourself to make career changes between labs and/or academia.

► **Entrepreneurship Opportunities & Skills**

This session will discuss some of the benefits, and drawbacks of pursuing entrepreneurial opportunities. It will describe what skills are critical for such ventures and how to identify entrepreneurial opportunities.

► **Publishing Your Research**

This session will discuss avenues for research publication, what is required for different types of publications, and ethical concerns of publishing such as plagiarism, dual submissions, and author ordering.

Plenary Sessions:

► **Finding Your Way: Overcoming Cultural Barriers**

This session will focus on cultural barriers that make pursuing an advanced degree in computing more challenging. Topics include family and culture, tools to help decide future when there is limited family support for research, familial pressure to provide economic support, and tradeoffs between money versus prestige.

► **Strategies for Human-Human Interaction**

This session will focus on strategies for productive interaction with colleagues (both faculty and students), including the opportunities and challenges of being an underrepresented minority and/or person with a disability in a computing technology career. Topics include interpersonal interaction dynamics, uncomfortable situations that might arise and how to react, recognizing and dealing with racial/cultural harassment and implicit bias, avoiding the pitfalls of being thought of as a token, and family-friendly policies to look for in a working environment.

Birds of a Feather:

► **Overcoming Insufficient Academic Preparation: Perceived and Real**

When striving to generate our best work, it matters that we have the necessary level of technical knowledge. How we and others perceive our academic preparation may impact our ability to leverage our technical expertise and to expand it as a project evolves.

This session discusses strategies to handle two scenarios:

- (1) A well-qualified student faces invalid assumptions of insufficient technical preparation, with minor gaps in technical expertise being addressed disrespectfully by project leaders or teammates;
- (2) A student notices inadequate background preparation and believes to be struggling much more than her/his peers. Without experience or guidance on how to catch up when behind in technical skills, the student may feel overwhelmed and end up unable to make progress.

► **Empowerment of People with Disabilities**

According to the Census Bureau's 2014 American Community Survey, 13% of the U.S. population has some disability; as shown in the National Science Foundation, National Center for Science and Engineering Statistics *Women, Minorities, and Persons with Disabilities in Science and Engineering Special Report* (<http://www.nsf.gov/statistics/wmpd/>), About 11% of undergraduate students reported a disability. Undergraduates with disabilities are older than those without disabilities and are slightly more likely to attend a 2-year institution. In addition, nearly one in four undergraduates with a disability enrolls in an STEM field. About 7% of graduate students reported a disability in 2012. About 20% of graduate students with disabilities are likely to enroll in an STEM field. We need to work, on a one-to-one basis, with people with disabilities from very young ages to build their self-esteem and feel comfortable to choose a STEM career. We need to work as well with their teachers and parents as, in many cases, they are discouraging young people with disabilities to embrace STEM careers.



Speakers

► Dorian Arnold



dorian.arnold@emory.edu **Emory University**

Dorian Arnold is an associate professor of Computer Science at Emory University. His research interests include operating and distributed systems, fault-tolerance, online (streaming) data analysis and high-performance tools. Particularly, he is interested in the scalability and reliability issues that abound in extreme scale computing environments comprising hundreds of thousands or even millions of components. Arnold's research group maintains strong collaborations with leading U.S. national labs and universities, including the Lawrence Livermore, Los Alamos and Sandia National Labs. In part due to such collaborations with world-class scientists and opportunities to work on cutting edge computing platforms, His research projects have won Top 100 R&D awards in 1999 and 2011. In 2017, he was selected as an ACM Distinguished Speaker. Arnold has held many leadership roles in major HPC conferences. Among other roles, he currently serves on the steering committee for the SC Conference and as an Associate Editor of the IEEE Transactions on Parallel and Distributed Systems. He is also very committed to diversity and inclusion in computer science most recently serving as General Chair for Tapia 2017. Arnold received Ph.D. and M.S. degrees in Computer Science from the Universities of Wisconsin and Tennessee, respectively. He also received his B.S. in Math and Computer Science from Regis University (Denver, CO) and his A.S. in Physics, Chemistry and Math from St. John's Junior College (Belize).

► Shiri Azenkot



shiri.azenkot@cornell.edu **Cornell Tech, Cornell University**

Shiri Azenkot is an Assistant Professor of Information Science at the Jacobs Technion-Cornell Institute at Cornell Tech, Cornell University. Her research interests are in accessibility and interaction on new platforms. Shiri frequently publishes at top HCI and accessibility conferences, including CHI, ASSETS, UIST, and UbiComp. Currently, her research is funded by the NSF, AOL, Verizon, and Facebook. Before arriving at Cornell Tech, she was a Ph.D. student in Computer Science & Engineering at the University of Washington, where she was advised by Richard Ladner and Jacob Wobbrock. Shiri has received the UW graduate medal (awarded to just one Ph.D. candidate at the university each year), a National Science Foundation Graduate Research Fellowship, and an AT&T Labs Graduate Fellowship.

► Alan Bivens



jbivens@us.ibm.com **IBM Research**

Alan Bivens is the Executive Assistant (Chief of Staff) for Marie Wieck, the General Manager of IBM Blockchain. In this role, Alan helps to drive strategy for the Blockchain Business Unit and related partners. Previously, Alan has held Senior Management roles in various parts of IBM Research leading global teams focused on bringing innovation to IBM offerings in the areas of AI, analytics, and blockchain. Alan received a Ph.D. in Computer Science from Rensselaer Polytechnic Institute in 2003 at which point he joined the Systems Research organization at the T.J. Watson Research Center. In his 15+ years at IBM, he has contributed to IBM products in Systems Software, Systems Hardware, Software Management, and Cloud and Support Services. Alan has invented/co-invented 40+ patents, authored/co-authored 40+ publications, and has led / contributed to several standards groups.

► Ramon Caceres



caceres@google.com Google

Ramon Caceres has more than 30 years of experience in technology research and development. His areas of focus include computer systems and networks, mobile computing systems and applications, and location data analysis and modeling. He is currently a Software Engineer at Google. He was previously a Lead Member of Technical Staff at AT&T Labs, a Research Staff Member at IBM Research, and Chief Scientist of Vindigo, a pioneering provider of location-based applications for mobile devices. He holds a Ph.D. in Computer Science from the University of California at Berkeley. He is an IEEE Fellow and an ACM Distinguished Scientist.

► Lori Clarke



CO-CHAIR, CRA-URMD GRADUATE COHORT WORKSHOP



clarke@cs.umass.edu University of Massachusetts Amherst

Lori A. Clarke is an emerita professor in the College of Information and Computer Sciences, University of Massachusetts Amherst, after serving on the computer science faculty for forty years and as chair from 2011-2015. She is a Fellow of the ACM and IEEE, and a board member of the Computing Research Association's Committee on the Status of Women in Computing Research (CRA-W). She is a former vice chair of the Computing Research Association (CRA), co-chair of CRA-W, IEEE Publication Board member, associate editor of ACM TOPLAS and IEEE TSE, member of the CCR NSF advisory board, and ACM SIGSOFT chair. Awards include the 2012 SIGSOFT Outstanding Research Award, 2011 University of Massachusetts Outstanding Accomplishments in Research and Creative Activity Award, the 2009 College of Natural Sciences and Mathematics Outstanding

Faculty Service Award, the 2004 University of Colorado, Boulder Distinguished Engineering Alumni Award, and the 2002 SIGSOFT Distinguished Service Award. Dr. Clarke's research is in the area of software engineering. She is one of the initial developers of symbolic execution and developed one of the first model checking systems applicable to software systems. She has also worked in requirements engineering and object management. Recently she has been investigating applying software engineering technologies to detect errors and vulnerabilities in complex, human-intensive processes in domains such as healthcare and digital government. She is also involved in efforts to increase participation of underrepresented groups in computing research.

► Jeanine Cook



jeacock@sandia.gov Sandia National Laboratories

Dr. Jeanine Cook is currently a Principal Member of Technical Staff at Sandia National Laboratories in the Scalable Architectures group. Her research interests include processing-in-memory architectures, next-generation memory technologies and subsystems, performance analysis tools for Exascale systems, and performance modeling and simulation. Prior to joining Sandia, Dr. Cook was an Associate Professor in The Klipsch School of Electrical and Computer Engineering at New Mexico State University. Dr. Cook was honored with a Presidential Early Career Award in Science and Engineering (PECASE) in 2008 from President George Bush for her work in performance modeling. She is a member of IEEE Computer Society and ACM.

► Dilma Da Silva



dilma@cse.tamu.edu Texas A&M University

Dilma Da Silva joined the Department of Computer Science and Engineering at Texas A&M University as its new department head in 2014. Her primary research interests are cloud computing, operating systems, distributed computing, and high-end computing. Before joining Texas A&M, she worked at Qualcomm Research in California (2012-2014), IBM Thomas J. Watson Research Center in New York (2000-2012) and the University of Sao Paulo in Brazil (1996-2000). Dilma is an ACM Distinguished Scientist, a member of the board of CRA-W (Computer Research Association's Committee on the Status of Women in Computing Research), a co-founder of the Latinas in Computing group, and an event liaison with USENIX. She served as an officer at ACM SIGOPS from 2011 to 2015 and chaired the ACM Senior Award Committee in 2015. She is an Associate Editor for several journals, has chaired 30 conferences/workshops and participated in more than 100 program committees. She has published more than 80 technical papers and filed 15 patents. Dilma received her doctoral degree in computer science from Georgia Tech in 1997 and her bachelor's and master's degrees from the University of São Paulo, Brazil.

► Jeff Forbes



forbes@cs.duke.edu **Duke University**

Jeff Forbes is an Associate Professor of the Practice of Computer Science at Duke University. He received his B.S. and Ph.D. Degrees in Computer Science from Stanford University and the University of California, Berkeley, respectively. His research interests include computer science education, social information processing, and learning analytics. Prof. Forbes is involved with a number of leading efforts to improve and broaden access to computer science education via organizations such as NSF, CMD-IT, ACM, CRA, and the CSforALL Consortium steering committee.

► Ayanna Howard



**CO-CHAIR, CRA-URMD GRADUATE
COHORT WORKSHOP**



ah260@gatech.edu **Georgia Institute of Technology**

Ayanna Howard, Ph.D., is Professor and Chair of the School of Interactive Computing in the College of Computing at the Georgia Institute of Technology. She also holds the Linda J. and Mark C. Smith Chair in Bioengineering. As an educator, researcher, and innovator, Dr. Howard's career focus is on intelligent technologies that must adapt to and function within a human-centered world. Her work, which encompasses advancements in artificial intelligence (AI), assistive technologies, and robotics, has resulted in over 200 peer-reviewed publications in a number of projects - from healthcare robots in the home to AI-powered STEM apps for children with diverse learning needs. She has over 20 years of R&D experience covering a number of projects that have been supported by various agencies including: National Science Foundation, NewSchools

Venture Fund, Procter and Gamble, NASA, and the Grammy Foundation. Dr. Howard received her B.S. in Engineering from Brown University, and her M.S. and Ph.D. in Electrical Engineering from the University of Southern California. To date, her unique accomplishments have been highlighted through a number of awards and articles, including highlights in USA Today, Upscale, and TIME Magazine, as well as being named a MIT Technology Review top young innovator. In 2013, she also founded Zyrobotics, which is currently licensing technology derived from her research and has released their first suite of STEM educational products to engage children of all abilities. From 1993-2005, Dr. Howard was at NASA's Jet Propulsion Laboratory, California Institute of Technology. She has also served as the Associate Director of Research for the Georgia Tech Institute for Robotics and Intelligent Machines and as Chair of the multidisciplinary Robotics Ph.D. program at Georgia Tech.

► Chad Jenkins



ocj@umich.edu **University of Michigan**

Odest Chadwicke Jenkins, Ph.D., is an Associate Professor of Computer Science and Engineering at the University of Michigan. Prof. Jenkins earned his B.S. in Computer Science and Mathematics at Alma College (1996), M.S. in Computer Science at Georgia Tech (1998), and Ph.D. in Computer Science at the University of Southern California (2003). He previously served on the faculty of Brown University in Computer Science (2004-15). His research addresses problems in interactive robotics and human-robot interaction, primarily focused on mobile manipulation, robot perception, and robot learning from demonstration. His research often intersects topics in computer vision, machine learning, and computer animation. Prof. Jenkins has been recognized as a Sloan Research Fellow in 2009. He is a recipient of the Presidential Early Career Award for Scientists

and Engineers (PECASE) for his work in physics-based human tracking from video. His work has also been supported by Young Investigator awards from the Office of Naval Research (ONR) for his research in learning dynamical primitives from human motion, the Air Force Office of Scientific Research (AFOSR) for his work in manifold learning and multi-robot coordination and the National Science Foundation (NSF) for robot learning from multivalued human demonstrations. Prof. Jenkins is a Senior Member of the ACM and the IEEE.

► Russ Joseph



rjoseph@eecs.northwestern.edu **Northwestern University**

Russ Joseph is an associate professor in the Electrical Engineering and Computer Science Department at Northwestern University (Evanston, IL). His primary research interest is in computer architecture, focusing on the design and implementation of power-aware and reliability-aware computer systems. He is the recipient of a 2007 NSF CAREER Award. He earned his Ph.D. in Electrical Engineering from Princeton University in 2004. He earned a B.S. degree with a double major in Electrical and Computer Engineering and Computer Science from Carnegie Mellon University in 1999.

► Shaun Kane



shaun.kane@colorado.edu **University of Colorado Boulder**

Shaun Kane is an Assistant Professor in the Department of Computer Science and, by courtesy, in the Department of Information Science at the University of Colorado, Boulder. He is also a CU ATLAS Faculty Fellow and director of the CU Superhuman Computing Lab. His research explores ways to make computing devices easier to use, especially for people with disabilities and people in distracting environments. His research has been supported by a Google Lime Scholarship, an NSF CAREER Award, and an Alfred P. Sloan Fellowship. He received his Ph.D. from The Information School at the University of Washington in 2011, and was an assistant professor in the Department of Information Systems at the University of Maryland, Baltimore County from 2011 to 2014.

► Daniel Lustig



dlustig@nvidia.com **NVIDIA**

Daniel Lustig is a Senior Research Scientist at NVIDIA. His research at NVIDIA focuses on computer architecture, memory consistency models, and memory system design. Academically, his “Check” line of work, which focuses on full-stack verification of memory consistency models from the programming language layer down to the microarchitecture, has been awarded three selections in IEEE Micro Top Picks. He also serves as the chair of the RISC-V Memory Consistency Model Task Group. Dan received his Ph.D. from Princeton in 2015 under the supervision of Margaret Martonosi. Prior to that, he received a MA from Princeton in 2011 and a BSE in Electrical Engineering and Mathematics from the University of Pennsylvania in 2009.

► Daniela Marghitu



marghda@auburn.edu **Auburn University**

Dr. Daniela Marghitu received her B.S. and M.S. degrees in Automation and Computing from Polytechnic University of Bucharest, and her Ph.D. degree in Automation and Computing from University of Craiova, European Community. Her teaching experience includes a variety of Information Technology and Computing courses (e.g., Introduction to Computing for Engineers and Scientists, Network Programming with HTML and Java, Personal Computer Applications, Spreadsheet-Based Applications with Visual BASIC, Web Application Development). Her research areas include the K12 Inclusive Computing Research and Outreach; Web Design; Education and Assistive Technology; Software Engineering; Web and Software Usability and Accessibility. Dr. Marghitu has participated in numerous administrative activities at Auburn University. Among these activities

are the following: Auburn University Board of Trustee Faculty Representative; Auburn University representative for National Center for Women in Information Technology; AccessComputing, AccessIOK, and AccessEngineering Alliances; technology coordinator for Alabama Alliance for Students with Disabilities in STEM Auburn University representative; Auburn University Persons with Disabilities Committee chair; Founder and Director Auburn University Laboratory for Education and Assistive Technology; faculty representative Auburn University Core Curriculum Oversight committee and Multicultural Diversity Commission; Chair World Usability Day Web Site Committee Chair; and Panel member for the National Science Foundation. Dr. Marghitu has received funding for research and education projects from National Science Foundation, National Center for Woman in Information Technology, Daniel F. Breeden Endowment for Faculty Enhancement, AccessComputing Alliance, Computer Science Collaboration Project, Microsoft Fuse Research, Altova Co., and Pearson Education Publishing Co. Dr. Marghitu published seven Information Technology

books at Pearson Publishing Co., articles at Transactions of the SDPS: Journal of Integrated Design and Process Science, User Experience Magazine, Journal of Computing Sciences in Colleges, International Journal for Virtual Reality, Journal of SMET Education and Research. She has published peer reviewed papers and gave presentations at numerous international conferences in USA, France, Germany, Spain, Portugal and Romania (e.g., ACM Special Interest Group on Computer Science Education Technical Symposium, International Technology and Persons with Disabilities Conference, International Conference on Software Engineering Advances, EDUCAUSE, Association for Advancement of Computing in Education, International Society for Technology in Education, Society for Design and Process Science, and American Society for Engineering Education). Among Dr. Marghitu's honors and awards are the following: 2011 AccessComputing Capacity Building Award, the 2012 Auburn University Access award, the 2012 Society for Design and Process Science Outstanding Achievement Award, the 2013 Microsoft Fuse Research award, the 2015 DO-IT Trailblazer award, the 2017 IARIA Fellowship, and the 2017 Society for Design and Process Science Outstanding Fellowship. She has been a member of the congressionally mandated NSF Committee on Equal Opportunities in Science and Engineering since 2015.

► Jason Mars



jason@clinc.com University of Michigan

Jason Mars is a professor of computer science at the University of Michigan where he directs Clarity Lab, one of the top labs in the world to be trained in AI and system design. He conducts collaborative research with IBM Watson, Google, Facebook, Intel and the National Science Foundation. Jason has devoted his career to solving difficult real-world problems with AI, computer vision, and natural language processing, at scale. Prior to The University of Michigan, Jason was a professor at UCSD. Jason's work building the world's first open-source platform for constructing large-scale AI and deep learning-based Intelligent Personal Assistants (IPAs) has been recognized globally and continues to have a significant impact on the tech industry and academia. Jason holds a Ph.D. in Computer Science from UVA. Jason has recently received prestigious awards including

ISCA Hall of Fame, The Google Research Award and The National Science Foundation's Career award.

► James Mickens



mickens@g.harvard.edu Harvard University

James Mickens is an associate professor of computer science at Harvard University. His research focuses on the performance, security, and robustness of large-scale web services. Mickens received a B.S. degree in computer science from the Georgia Institute of Technology, and a Ph.D. in computer science from the University of Michigan. Before coming to Harvard, he spent six years as a researcher at Microsoft.

► Jaime Moreno



jhmoreno@us.ibm.com IBM Research

Dr. Jaime H. Moreno is Distinguished Researcher, Senior Manager, Data Centric High Performance Systems, at the IBM Thomas J. Watson Research Center in New York, where he focuses on next-generation supercomputing systems. He joined the IBM Research Division in 1992, and has led various teams on microprocessor and high-performance system architecture, design and performance analysis. He has published multiple papers, is coauthor of two books, including a textbook translated to Portuguese and Chinese, holds many patents in processor architecture, and has been recognized as Master Inventor at IBM Research. He has been actively involved with the research community through organization and participation in major conferences in his field. His current research interests address future supercomputers systems,

applications and their co-design, and the convergence of Cognitive Computing with High-Performance Computing. Before joining IBM, he was a faculty member at the University of Concepcion, Chile. He received his Ph.D. and M.S. degrees in Computer Science from the University of California Los Angeles, and a degree in Electrical Engineering from the University of Concepcion, Chile. He is member of IEEE and ACM.

► Melanie Moses



melaniem@unm.edu **University of New Mexico**

Professor Melanie Moses earned a B.S. from Stanford University in Symbolic Systems, an interdisciplinary program in cognition and computation, and a Ph.D. in Biology from the University of New Mexico in 2005. She is currently an Associate Professor in the Department of Computer Science at the University of New Mexico and External Faculty at the Santa Fe Institute. Her interdisciplinary research exists at the boundaries of Computer Science and Biology with over 50 peer reviewed publications in computational and mathematical biology and biologically-inspired swarm robotics. Research in the Moses Lab focuses on computational modeling of complex biological systems, particularly on cooperative search strategies in immune systems and ant colonies. Her research also applies principles from biology to design computational systems, particularly robotic swarms that replicate ant behaviors to perform collective tasks. Her research lab includes 14 includes post docs, undergraduate and graduate students and high school interns from Computer Science and Biology. Professor Moses was the co-director of the NIH funded UNM Program in Interdisciplinary Biological and Biomedical Sciences 2013 - 2015, and directs the CSforAll course, an introductory programming course in computer modeling and simulation in which 400 New Mexico high school students have been introduced to computer science and earned dual credit at UNM. Professor Moses is the Principal Investigator for the NASA Swarmathon, a swarm robotics competition that aims to engage 1000 students from Minority Serving Institutions to develop new swarm robotic algorithms to revolutionize space exploration. She is honored to have been a Ford Foundation Dissertation Diversity Fellow and a Microsoft Research New Faculty Fellowship Finalist, and to have received the UNM Outstanding New Teacher of the Year Award and the School of Engineering New Faculty Awards for Excellence in Teaching and Research.

► Kunle Olukotun



CO-CHAIR, CRA-URMD GRADUATE COHORT WORKSHOP



kunle@Stanford.EDU **Stanford University**

Kunle Olukotun is the Cadence Design Systems Professor of Electrical Engineering and Computer Science at Stanford University. Olukotun is well known as a pioneer in multicore processor design and the leader of the Stanford Hydra chip multiprocessor (CMP) research project. Olukotun founded Afara Websystems to develop high-throughput, low-power multicore processors for server systems. The Afara multicore processor, called Niagara, was acquired by Sun Microsystems. Niagara derived processors now power all Oracle SPARC-based servers. Olukotun currently directs the Stanford Pervasive Parallelism Lab (PPL), which seeks to proliferate the use of heterogeneous parallelism in all application areas using Domain Specific Languages (DSLs). Olukotun is a member of the Data Analytics for What's Next (DAWN) Lab which is developing infrastructure for usable machine learning. Olukotun is an ACM Fellow and IEEE Fellow for contributions to multiprocessors on a chip and multi-threaded processor design. Olukotun received his Ph.D. in Computer Engineering from The University of Michigan.

► Mary Lou Soffa



CO-CHAIR, CRA-URMD GRADUATE COHORT WORKSHOP



soffa@virginia.edu **University of Virginia**

Mary Lou Soffa is the Owen R. Cheatham Professor of Sciences in the Computer Science Department at the University of Virginia, serving as the Department Chair from 2004 to 2012. From 1977 to 2004, she was a Professor of Computer Science at the University of Pittsburgh and also served as the Dean of Graduate Studies in the College of Arts and Sciences for five years. Her research interests include cloud computing, warehouse scale computers, software systems for multi-core architectures, optimizing compilers, software testing, and program analysis. She has directed 32 Ph.D. students to completion, half of whom are women. Mary Lou is a Fellow of the Association for Computing Machinery (ACM) and a Fellow of The Institute of Electrical and Electronic Engineers (IEEE). In 2012, she received the Ken Kennedy Award for contributions to compiler technology and software engineering, exemplary service to the profession, and lifelong dedication to mentoring and improving diversity in computing. She received the Anita Borg Technical Leadership Award in 2011, which celebrates a woman who led or developed a product, process, or innovation that made a notable impact on business or society. Other awards she received include the ACM SIGSOFT Influential Educator Award in 2014, the Distinguished Alumni Award from the Computer Science Department at the University of Pittsburgh in 2017, the Computing Research Association (CRA) A. Nico Habermann Award in 2006, a Girl Scout Woman of Distinction Award in 2003, and the White House's Presidential Award for Excellence in Science, Mathematics and Engineering Mentoring in 1999. She has also received a number of distinguished paper awards. Mary Lou was a member of the CRA Board for ten years. She served as a co-chair and member on the CRA-W Board and co-founded the CRA-W Graduate Cohort

Program and the CRA-W Mentoring Program for Associate Professors. Mary Lou served on ACM Council from 2000-2016. She has served on the editorial board of a number of journals, including the *ACM Transactions on Programming Languages and Systems* and the *IEEE Transactions on Software Engineering*. In addition, she has been a conference chair, program chair or program committee member for many conferences in programming languages, software systems and software engineering.

► Armando Solar-Lezama



asolar@csail.mit.edu MIT

Associate Professor Armando Solar-Lezama leads the Computer Aided Programming group at MIT and works at the intersection of Programming Systems and Artificial Intelligence. His is best known for his work on program synthesis, and the early development of the Sketch synthesis system. He was also an early proponent of the use of program execution data as a way to guide the synthesis process and of the use of synthesis to enable programming through very high-level interactions such as whiteboard-style diagrams of data-structure manipulations. His work has been applied in a number of domains, from databases and networking to high-performance scientific computing.

► Lydia Tapia



tapia@cs.unm.edu University of New Mexico

Lydia Tapia is an Associate Professor in the Department of Computer Science at the University of New Mexico. She received her Ph.D. in Computer Science from Texas A&M University and her B.S. in Computer Science from Tulane University. Her research contributions are in the development of computationally efficient algorithms for the simulation and analysis of high-dimensional motions for robots and molecules. Specifically, she explores problems in computational structural biology, motion under stochastic uncertainty, and reinforcement learning. Based on this work, she has been awarded two patents, one on a novel unmanned aerial vehicle design and another on a method to design allergen treatments. Lydia is the recipient of the 2016 Denice Denton Emerging Leader ABIE Award from the Anita Borg Institute, a 2016 NSF CAREER Award for

her work on simulating molecular assembly, and the 2017 Computing Research Association Committee on the Status of Women in Computing Research (CRA-W) Borg Early Career Award. When not in the lab, you can often find Lydia and her students doing interactive robot demos for the community at local schools, robotics competitions, and museums.

► C.J. Taylor



cjtaylor@cis.upenn.edu University of Pennsylvania

Dr. Taylor received his A.B. degree in Electrical Computer and Systems Engineering from Harvard College in 1988 and his M.S. and Ph.D. degrees from Yale University in 1990 and 1994 respectively. Dr. Taylor was the Jamaica Scholar in 1984, a member of the Harvard chapter of Phi Beta Kappa and held a Harvard College Scholarship from 1986-1988. From 1994 to 1997, Dr. Taylor was a postdoctoral researcher and lecturer with the Department of Electrical Engineering and Computer Science at the University of California, Berkeley. He joined the faculty of the Computer and Information Science Department at the University of Pennsylvania in September 1997. He received an NSF CAREER award in 1998 and the Lindback Minority Junior Faculty Award in 2001. In 2012 he received a best paper award at the IEEE Workshop on the Applications of Computer Vision. Dr

Taylor's research interests lie primarily in the fields of Computer Vision and Robotics and include: reconstruction of 3D models from images, vision-guided robot navigation and smart camera networks. Dr. Taylor has served as an Associate Editor of the *IEEE Transactions of Pattern Analysis and Machine Intelligence*. He has also served on numerous conference organizing committees and was a Program Chair of the 2006 and 2017 editions of the IEEE Conference on Computer Vision and Pattern Recognition and of the 2013 edition of 3DV. In 2012 he was awarded the Christian R. and Mary F. Lindback Foundation Award for Distinguished Teaching at the University of Pennsylvania.

Institutions

We're delighted to host 100 participants representing over 60 institutions across the United States and Canada this year.

Boise State University ♦ Clemson University ♦ Columbia University ♦ Florida State University ♦ George Washington University ♦ Georgia Institute of Technology ♦ Georgia State University ♦ Howard University ♦ Indiana University (System) ♦ Iowa State University of Science and Technology ♦ Jackson State University ♦ Kansas State University ♦ Michigan Technological University ♦ Mississippi State University ♦ Missouri University of Science and Technology ♦ Montana State University ♦ New Mexico State University ♦ North Carolina Agricultural and Technical State University ♦ North Carolina State University ♦ Northeastern University ♦ Northwestern University ♦ Nova Southeastern University ♦ Princeton University ♦ Rice University ♦ Rochester Institute of Technology ♦ Rutgers, The State University of New Jersey ♦ Southern New Hampshire University ♦ State University of New York at Buffalo ♦ Syracuse University ♦ Texas A&M University ♦ The Tulane University of New Orleans ♦ The University of Oklahoma ♦ Tufts University ♦ Universidad Politecnica de Puerto Rico ♦ University of Baltimore ♦ University of California, Berkeley ♦ University of California, Riverside ♦ University of Central Florida ♦ University of Cincinnati ♦ University of Florida ♦ University of Illinois at Urbana-Champaign ♦ University of Kentucky ♦ University of Lethbridge ♦ University of Maryland at College Park ♦ University of Maryland Baltimore County ♦ University of Maryland University College ♦ University of Massachusetts at Amherst ♦ University of Memphis ♦ University of Michigan - Ann Arbor ♦ University of New Orleans ♦ University of North Texas ♦ University of Notre Dame ♦ University of Saskatchewan ♦ University of Southern California ♦ University of Victoria ♦ University of Washington ♦ University of Waterloo ♦ Virginia Polytechnic Institute and State University ♦ Washington University in St. Louis ♦ Wayne State University ♦ Worcester Polytechnic Institute

Sponsors





Poster Session

#	Primary Research Area	Name	Last	Current University or Institution	Abstract Title
1	Artificial Intelligence	Daniel	Seita	University of California, Berkeley	An Efficient Minibatch Acceptance Test for Metropolis-Hastings
2	Bioinformatics	Sampson	Akwafuo	University of North Texas	Investigating Impacts of HIV Epidemic Interventions Among Key Population in Nigeria: A Modelling Approach
3	Bioinformatics	Anna	Kirkpatrick	Georgia Institute of Technology	RNA Pairing Properties under Uniform And Thermodynamically-Weighted Distributions
4	Bioinformatics	Paz	Zait-Givon	Rice University	Processing Intracardiac Data to Map Conduction Pathways
5	Computer Architecture	Cesar	Gomes	Tufts University	Capacity Curves - quantifying workload utility
6	HCI	Denae	Ford	North Carolina State University	"We Don't Do That Here": How Collaborative Editing with Mentors Improves Engagement in Social Q&A Communities
7	HCI	Noella	Kolash	Rochester Institute of Technology	MarsU: Usability Study of a Digital Game Based Learning Tool for Deaf Learners
8	HCI	Juan	Maestre	Indiana University at Bloomington	The Suitability of the ARC Method for Studying Stigmatized Populations
9	Information Retrieval	Oghenemaro	Anuyah	Boise State University	Extending Safe Search Functionality for Identifying Child-friendly Web Resources
10	Machine Learning	Francisco	Garcia	University of Massachusetts, Amherst	Leveraging External Knowledge for Better Decision Making
11	Machine Learning	Tomas	Geffner	University of Massachusetts, Amherst	Efficient Variational Inference
12	Machine Learning	Irish	Medina	University of Waterloo	Forecasting Short Term Water Consumption for Multi-Family Residences
13	Natural Language Processing	Elozino	Egonmwan	University of Lethbridge	Abstractive Multi-document Text Summarization-Fusing, Paraphrasing and Compressing with Deep Neural Networks
14	Regression Analysis	Samuel	Anyaso-Samuel	Boise State University	Some Contributions to the Interpretation of Fuzzy Regression Intervals
15	Security/Privacy	Matthew	Landen	Georgia Institute of Technology	Android Malware Detection using event sequences
16	Social Computing	Bryan	Dosono	Syracuse University	Identity Work on Reddit: Toward Social Support and Collective Action of Asian Americans and Pacific Islanders
17	Software Engineering	Miguel	Jimenez Achinte	University of Victoria	Deployment Specification Challenges in the Context of Large-Scale Systems
18	Software Engineering	Wing	Lam	University of Illinois, Urbana-Champaign	A characteristic study of bug reports



Poster Abstracts

■ Poster #1 • Daniel Seita

University of California, Berkeley

Artificial Intelligence | An Efficient Minibatch Acceptance Test for Metropolis-Hastings

We present a novel Metropolis-Hastings method for large datasets that uses small expected-size minibatches of data. Previous work on reducing the cost of Metropolis-Hastings tests yield variable data consumed per sample, with only constant factor reductions versus using the full dataset for each sample. Here we present a method that can be tuned to provide arbitrarily small batch sizes, by adjusting either proposal step size or temperature. Our test uses the noise-tolerant Barker acceptance test with a novel additive correction variable. The resulting test has similar cost to a normal SGD update. Our experiments demonstrate several order-of-magnitude speedups over previous work. We additionally show that our test can be used with neural networks.

■ Poster #2 • Sampson Akwafuo

University of North Texas

Bioinformatics | Investigating Impacts of HIV Epidemic Interventions Among Key Population in Nigeria: A Modelling Approach

Using Females Who Sell Sex (FWSS) in Nigeria as a case study, this research develops a novel risk equation for estimating new infections among FWSS, their clients and communities. It uses a hybrid SUDT and SIT structural model. It considers number of contacts, number of protected and unprotected sexual acts, population and other existing values as base inputs. Using python programming, the impacts of FWSS direct interventions on their clients, their female partners and the general population is estimated. The levels of the program implementation, needed on each scenario, to achieve the required number of averted new infections are also modelled. Our model can be used to estimate the risk of a population set to a sexually transmitted disease. Public health workers can use the model to prepare a fit-for-purpose intervention program for specific community members.

■ Poster #3 • Anna Kirkpatrick

Georgia Institute of Technology

Bioinformatics | RNA Pairing Properties under Uniform And Thermodynamically-Weighted Distributions

Understanding the structure of RNA is a problem of significant interest to biochemists. An RNA molecule has both a linear structure, consisting of a sequence of nucleotides, and secondary and tertiary structures determined by how these nucleotides interact. If the linear structure is known, algorithmically determining even the secondary structure, which can be viewed as a non-crossing matching, is a significant challenge, in part because the thermodynamic energy function used to make such predictions does not perform well on longer sequences. This work attempts to understand the relationship between this energy function and properties of the structures it predicts. Markov chain Monte Carlo computations are used to determine the distribution of certain properties under the thermodynamic energy function, and these are compared to properties of structures sampled uniformly. Some relevant combinatorial results on non-crossing perfect matchings are also presented.

■ Poster #4 • Paz Zait-Givon Rice University

Bioinformatics | Processing Intracardiac Data to Map Conduction Pathways

Tracing activation intervals along conduction pathways through the heart is a fundamental step in understanding the progression of heart failure and developing new and personalized therapies. This work will utilize intracardiac data from several locations along the heart and well-known signal processing techniques to delimitate activation intervals across the heart and monitor the activation intervals and delays.

■ Poster #5 • Cesar Gomes Tufts University

Computer Architecture | Capacity Curves - quantifying workload utility

Last level cache management has seen many iterations and strategies in recent years. Cache replacement, prefetching, coherence, and partitioning are the big categories in which academics and industry seek to improve last level cache efficiency. Inefficiencies in the usage of large last level caches are exacerbated by increasing core counts and data intensive workloads. Further, increasingly complex techniques exhibit diminished or non-existent improvements on a per workload basis as cache size increases. We present Cache Capacity Curves to represent workload fit within common last level cache sizes. In this work, we use Cache Capacity Curves to classify workloads. We formalize a classification method based on these “curves,” and determine the impact cache management techniques have on classification (e.g. predictor, pre-fetcher, etc.). Further, a case study on way partitioning is conducted to demonstrate how a given last level cache should share resources between two workloads given knowledge of respective curves.

■ Poster #6 • Denae Ford North Carolina State University

HCI | “We Don’t Do That Here”: How Collaborative Editing with Mentors Improves Engagement in Social Q&A Communities

Online question-and-answer (Q&A) communities like Stack Overflow have norms that are not obvious to novice users. Novices create and post programming questions without feedback, and the community enforces site norms through public downvoting and commenting. This can leave novices discouraged from further participation. We deployed a month long, just-in-time mentorship program to Stack Overflow in which we redirected novices in the process of asking a question to an on-site Help Room. There, novices received feedback on their question drafts from experienced Stack Overflow mentors. We present examples and discussion of various question improvements including: question context, code formatting, and wording that adheres to on-site cultural norms. We find that mentored questions are substantially improved over non-mentored questions, with average scores increasing by 50%. We provide lessons learned and implications that challenge how socio-technical communities design onboarding experiences across domains.

■ **Poster #7 • Noella Kolash**
Rochester Institute of Technology

HCI | MarsU: Usability Study of a Digital Game Based Learning Tool for Deaf Learners

Deaf American Sign Language (ASL) users often learn English as a second language (L2). L2 learners often experience difficulties learning statistics. A Digital Game-Based Learning tool was developed—viz. MarsU—to aid the learning of statistics. MarsU was administered to 6 deaf students for a usability study. A pre- and post-test was administered to measure learning along with the Test of Silent Contextual Reading Fluency, Motivated Strategies for Learning Questionnaire, System Usability Scale, and Net Promoter. The entire session was recorded using two screen recorders and two web cameras to capture interaction with the game and task of instructions. Preliminary results show most participants did not understand the purpose of the game. The lengthy text-based tutorial was not helpful and participants could not recall the information in the tutorial. Several participants did not collect samples, which would inform their decisions toward winning the game. Preliminary results suggest the development of tutorials specifically designed for deaf learners is needed.

■ **Poster #8 • Juan Maestre**
Indiana University at Bloomington

HCI | The Suitability of the ARC Method for Studying Stigmatized Populations

Researchers in HCI have typically relied on face to face (FtF) methods for recruitment and data collection in their research with people living with HIV, whereas social scientists have adopted computer-mediated approaches to address concerns about data validity and access to this stigmatized population. We used the asynchronous remote community (ARC) research method to leverage HCI instruments in an online format. ARC successfully engaged people living with HIV by providing a safe space to discuss their experiences. By expanding on past ARC studies, we contribute to an ongoing conversation about defining ARC and working towards increased data validity – especially in stigmatized communities.

■ **Poster #9 • Oghenemaro Anuyah**
Boise State University

Information Retrieval | Extending Safe Search Functionality for Identifying Child-friendly Web Resources

One of the largest community of web users utilizing search engines are children. Children commonly use search engines to access sites that interest them such as social networking and online gaming, as well as for school work. While the use of search engines can facilitate children's access to web resources that address their information need, children may access resources considered inappropriate for them [31]. To prevent children's access to inappropriate resources, Google and Yahoo have made available a safe search option [3], and introduced Kidrex and Yahoooligans respectively, as children-oriented search engines that adopts safe search functionality. However, there are some cases where safe search strictly filters web resources. In this paper, we design a novel approach for filtering inappropriate resources from the retrieved set, without disregarding those ones that are relevant to the educational context of a child. We train a gradient boosting tree classifier with child-friendly and non-child friendly documents.

■ **Poster #10 • Francisco Garcia** **University of Massachusetts, Amherst**

Machine Learning | Leveraging External Knowledge for Better Decision Making

The area of reinforcement learning (RL) has seen a surge in interest in recent years. At the core of the RL problem lays that of exploration/exploitation, that is, when should a learning agent act with the goal of gathering more information about the environment (exploration) and when should it act according to what it has learned thus far (exploitation). Current research mostly focuses on the objective of obtaining well-performing policies, but largely ignores how an agent should behave while exploring. If we assume the agent has already experienced similar environments, it should be able to explore using prior knowledge. In this work, we propose a framework for directly optimizing how an agent should behave during exploration. We define a parametric exploration policy and optimize a user-defined notion of what it means to “explore well”.

■ **Poster #11 • Tomas Geffner** **University of Massachusetts, Amherst**

Machine Learning | Efficient Variational Inference

The use of probabilistic models comes with many advantages. Allowing experts in different areas to encode domain knowledge in a simple way being one of these. This makes it possible to train models with a low number of samples, which is extremely useful in situations where a limited number of data points are available. Also, probabilistic models have been used to infer structures that can be used to explain observations, understand the data better and make predictions about new samples. However, exact inference on these models is often intractable. Different approximate methods were developed, such as Variational Inference (VI). VI transforms the inference problem into an optimization one, which is often solved using Stochastic Gradient Descent. Situations arise in which high variance gradient estimates are obtained, complicating the optimization problem. In this work I propose a new way of obtaining control variates to reduce the variance of these gradient estimates.

■ **Poster #12 • Irish Medina** **University of Waterloo**

Machine Learning | Forecasting Short Term Water Consumption for Multi-Family Residences

Smart water meters have been installed across Abbotsford, British Columbia, Canada, to measure the hourly water consumption of consumers in the area. Using this water consumption data, we aim to develop machine learning models to predict water consumption in the short term for multi-family residences in the city of Abbotsford. Our work addresses the gap in research for predicting water consumption for multi-family residences and accounts for the current trend towards denser living spaces in urban areas.

■ **Poster #13 • Elozino Egonmwan** **University of Lethbridge**

Natural Language Processing | Abstractive Multi-document Text Summarization- Fusing, Paraphrasing and Compressing with Deep Neural Networks

Researchers, media analysts, meteorologists and humans in general, often need to access, digest and condense large amounts of data into informative highlights, as either abstracts, news digests, weather forecasts or other rich summarized forms. Hence, the need for automatic text summarization continues to be an important problem. We propose a multi-document abstractive text summarization system using deep neural networks for fusing, paraphrasing and compressing relevant sentences of the source texts. The conjuncture is that a good abstractive model should involve a combination of these three operations: sentence fusion, paraphrasing and compression. The proposed model, aims to fill the increasing demand for high-quality summaries that are grammatically correct, abstractive, coherent, informative, and non-redundant.

■ **Poster #14 • Samuel Anyaso-Samuel**
Boise State University

Regression Analysis | Some Contributions to the Interpretation of Fuzzy Regression Intervals

Fuzzy regression interval has been very useful in modelling and forecasting problems. In this work, we introduce the use of valid interpretation of the fuzzy system model with fuzzy regression approach such as conceptualizing fuzzy regression using OLS and minimum fuzziness criterion. The results shows a clear implementation and interpretation of fuzzy regression using the fuzzy regression interval.

■ **Poster #15 • Matthew Landen**
Georgia Institute of Technology

Security/Privacy | Android Malware Detection using event sequences

The Google Play Store serves as the app market for millions of android device users. This large userbase has become a target for malicious developers to distribute their dangerous applications. This research attempts to distinguish malicious apps from benign apps by examining the sequences of actions that invoke sensitive Android APIs. Leveraging a modified version of Intellidroid we extract the event sequences and feed them into our classification pipeline. The classification procedure consists of two steps. First, a long short term memory network is used to perform event sequence classification. Then, a support vector machine uses the network's results to predict whether the app is malicious or not. This process achieves an accuracy of greater than 90% on a sizable collection of apps. This method is useful or app stores to prevent dangerous applications from being published as well as for android users to better protect their phone and sensitive information.

■ **Poster #16 • Bryan Dosono**
Syracuse University

Social Computing | Identity Work on Reddit: Toward Social Support and Collective Action of Asian Americans and Pacific Islanders

Asian Americans and Pacific Islanders (AAPIs) are perceived as the "model minority" with a monolithic identity, in contrast to other marginalized racial groups in the United States. In reality, they are composed of different ethnicities, socio-economic backgrounds, and political ideologies. My research employs social network analysis with qualitative research methods to explore, interpret, and visualize large collections of social media data. I seek to understand how Asian Americans and Pacific Islanders (AAPIs) construct and express their identity in online communities and my dissertation research uncovers the ways in which AAPIs negotiate collective action in the context of online identity work.

■ **Poster #17 • Miguel Jimenez Achinte**
University of Victoria

Software Engineering | Deployment Specification Challenges in the Context of Large-Scale Systems

Traditionally, the focus of software deployment has been mainly on the infrastructure to realise deployment and configuration (D&C) of complex and distributed systems, with an increasing interest in deployment of internet of things and cyber-physical systems. Advances in job scheduling, storage orchestration, containerized applications, along with agile practices such as continuous integration and microservices architecture, have improved the state of the practice. However, little effort has been devoted to the need for D&C specifications to support the various levels of detail and abstraction present in large-scale systems. The understanding of the software components hierarchy has shifted from the comprehension of design artefacts, usually specified with static diagrams, to the understanding of runtime concepts. The DevOps movement has dramatically influenced how and when deployment is realized, but little has been done from the software perspective in terms of documentation and linkage between design and runtime artefacts in the sense of software specification as such. In this poster, I will present our motivation to focus on specification, along with a set of deployment specification requirements and challenges.

■ **Poster #18 • Wing Lam**
University of Illinois, Urbana-Champaign

Software Engineering | A characteristic study of bug reports

Developers often use a bug tracking system to track the bugs that appear in their software. Users of the software can utilize the system to report bugs that they encounter when using the software. Since the familiarity and expertise of the user to the software can greatly vary between users, the quality of the bug reports produced by users also greatly varies. Our work presents a characteristic study of the bug reports from popular open source projects. We study the bug reports that users wrote for these open source projects and the patches and test cases written by developers for such bug reports. Our study findings provide valuable insights into the common pitfalls of bug reports and the feasibility of them to be used manually and automatically to generate patches and test cases.



The Computing Research Association is dedicated to increasing the success and participation of women, underrepresented minorities and persons with disabilities in computing research. Diversity Drives Innovation!

We would like to thank the CRA URMD Grad Cohort Co-Chairs and Steering Committee for their assistance in organizing this workshop:

Co-Chairs:

Lori Clarke, University of Massachusetts Amherst
Ayanna Howard, Georgia Institute of Technology

Kunle Olukotun, Stanford University
Mary Lou Soffa, University of Virginia

Steering Committee:

Nancy Amato, Texas A&M University
Theo Benson, Duke University
Dilma Da Silva, Texas A&M University
Jeff Forbes, Duke University
Julia Hirschberg, Columbia University
Chad Jenkins, University of Michigan
Daniel Jimenez, Texas A&M University

Russ Joseph, Northwestern University
Richard Ladner, University of Washington
Jason Mars, University of Michigan
James Mickens, Harvard University
Melanie Moses, University of New Mexico
Armando Solar-Lezama, Massachusetts Institute of Technology
Lydia Tapia, University of New Mexico
Valerie Taylor, Argonne National Laboratory

www.cra.org

@CRAtweets

#CRAGradCohort

A special thanks to CRA staff for their assistance in organizing this event:

Erik Russell

Claire Brady • Daniela Cardenas

Sandra Corbett • Sabrina Jacob • Shar Steed

We would like to thank the following organizations for their generous contributions:

PLATINUM SPONSORS



SILVER SPONSORS



ASSOCIATION FOR COMPUTING MACHINERY

BRONZE SPONSORS



Wireless Code:

U: CRAGradCohort

P: 2018CRA