AYLIN CALISKAN

Assistant Professor
Information School
University of Washington

ACADEMIC APPOINTMENTS

- Assistant Professor, *September 2021 - present*
  University of Washington, Information School

- Assistant Professor, *July 2018 - August 2021*
  George Washington University, Department of Computer Science

- Core Faculty, *July 2019 - August 2021*
  Institute for Data, Democracy & Politics (funded by the Knight Foundation)
  George Washington University, School of Media and Public Affairs

- Postdoctoral Research Associate and CITP Fellow, *September 2015 - June 2018*
  Princeton University, Center for Information Technology Policy

    Advisor: Arvind Narayanan
    Research in bias and ethics in AI, privacy and security via machine learning and natural language processing

EDUCATION

- Ph.D. in Computer Science, *June 2015*
  Drexel University, College of Computing & Informatics

    Advisor: Rachel Greenstadt
    Dissertation: Stylometric Fingerprints and Privacy Behavior in Textual Data
    Research in privacy and security via machine learning and natural language processing

    Visiting Scholar, *June 2014 - August 2014*
    United States Army Research Laboratory, Adelphi, MD
    Research on source code authorship attribution

- M.S. in Computer Science, *June 2013*
  Drexel University, College of Computing & Informatics

- M.S.E. in Robotics, *May 2011*
  University of Pennsylvania, School of Engineering and Applied Sciences
  Thesis: Perception and Grasping in a Kitchen Environment
  Research in machine learning, computer vision, and localization

- B.S. in Information Systems, *May 2009*
  SUNY Binghamton University, Thomas J. Watson School of Engineering and Applied Sciences

- B.S. in Information Systems Engineering, *May 2009*
  Istanbul Technical University, Faculty of Electrical and Electronics Engineering

- Robert College
  High School Degree in Science and Mathematics, *June 2004*

HONORS AND AWARDS

- Invited Keynote. Implications of Biased AI on Democracy, Equity, and Justice. Natural Language Processing for Internet Freedom - COLING Workshop (2020)

- Speaker, National Academy of Sciences. Authenticity, Integrity, and Security in a Digital World (2019)

- Selected Rising Star in EECS (Electrical Engineering and Computer Science) at Stanford University (2017)

- Invited Keynote. Beyond Big Data: What Can We Learn from AI Models? AISec - CCS Workshop (2017)

- Best Talk Award, A Story of Discrimination and Unfairness: Implicit Bias Embedded in Language Models, HotPETS (2016)

- Invited Keynote. Natural Language Processing and Privacy: A Double Edged Sword. Infer - PETS Workshop (2016)
- Teaching Assistant Excellence Award Nominee, Graduate-Level Machine Learning Course, Drexel University (2015)
- IEEE Symposium on Security and Privacy Travel Award (2014)
- Inaugural College of Computing & Informatics Day PhD Poster Award, Doppelgänger Finder: Taking Stylometry To The Underground, Drexel University (2014)
- Upsilon Pi Epsilon Executive Council Award (2013)
- Brocade Scholarship, Grace Hopper Celebration of Women in Computing (2013)
- Xerox Scholarship, Grace Hopper Celebration of Women in Computing (2012)
- Andreas Pfitzmann Best Student Paper Award 2012, PETS (2012)
- Upsilon Pi Epsilon, International Honor Society in the Computing and Information Disciplines (2012)
- Dean's Fellowship, Drexel University College of Engineering: (2011, 2012, 2013, 2014)

MANUSCRIPTS

1. Social Biases in Word Embeddings and Their Relation to Human Cognition
   Aylin Caliskan and Molly Lewis
   The Atlas of Language Analysis in Psychology, Guilford Press 2020. Editors Morteza Deghani and Ryan Boyd.
   - book chapter preprint, 2020.
2. Send a Raven: Well-Behaved Covert Channels Seldom Make History
   Ryan Wails, Andrew Stange, Samanta Troper, **Aylin Caliskan**, Roger Dingledine, Rob Jansen, Micah Sherr
   - under submission, 2021.

PEER-REVIEWED PUBLICATIONS

3. ValNorm Quantifies Semantics to Reveal Consistent Valence Biases Across Languages and Over Centuries
   Autumn Toney and **Aylin Caliskan**
   Empirical Methods in Natural Language Processing (EMNLP 2021)
4. Disparate Impact of Artificial Intelligence Bias in Ridehailing Economy's Price Discrimination Algorithms
   Akshat Pandey and **Aylin Caliskan**
   AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AAAI/ACM AIES 2021) (Acceptance rate: 27%)
5. Detecting Emergent Intersectional Biases: Contextualized Word Embeddings Contain a Distribution of Human-like Biases
   Wei Guo and **Aylin Caliskan**
   AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AAAI/ACM AIES 2021) (Acceptance rate: 9%)
6. Image Representations Learned With Unsupervised Pre-Training Contain Human-like Biases
   Ryan Steed and **Aylin Caliskan**
   The 2021 ACM Conference on Fairness, Accountability, and Transparency (ACM FAccT), 2021. (Acceptance rate: 25%)
7. A Set of Maximally Distinct Facial Traits Learned by Machines is not Predictive of Appearance Bias in the Wild
   Ryan Steed and **Aylin Caliskan**
   AI and Ethics, 2021.
8. Automatically Characterizing Targeted Information Operations Through Biases Present in Discourse on Twitter
   Autumn Toney, Akshat Pandey, Wei Guo, David Broniatowski, and **Aylin Caliskan**
   15th IEEE International Conference on Semantic Computing (ICSC), 2021. (Acceptance rate: 20%)
9. If I Tap It, Will They Come? An Introductory Analysis of Fairness in a Large-Scale Ride Hailing Dataset
   Aylin Caliskan and Begum Kaplan
   Academy of Marketing Science Annual Conference (AMS), 2020.
10. Git Blame Who?: Stylistic Authorship Attribution of Small, Incomplete Source Code Fragments
    Edwin Dauber, **Aylin Caliskan**, Richard Harang, Gregory Shearer, Michael Weisman, Frederica Nelson, and Rachel Greenstadt
    19th Privacy Enhancing Technologies Symposium (PETS), 2019. (Acceptance rate: 16%)
11. When Coding Style Survives Compilation: De-anonymizing Programmers from Executable Binaries
    Aylin Caliskan, Fabian Yamaguchi, Edwin Dauber, Richard Harang, Konrad Rieck, Rachel Greenstadt, and Arvind Narayanan
    Network and Distributed System Security Symposium (NDSS), 2018. (Acceptance rate: 15%)
12. Semantics derived automatically from language corpora contain human-like biases
    Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan
    ***Science*, 2017. (Impact factor: 37.2%, cited 1,368 times)**

13. A Story of Discrimination and Unfairness
    Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan
    9th Hot Topics in Privacy Enhancing Technologies (HotPETs), 2016.
    Best Talk Award, HotPETS 2016

14. De-anonymizing Programmers via Code Stylometry
    Aylin Caliskan-Islam, Richard Harang, Andrew Liu, Arvind Narayanan, Clare Voss, Fabian Yamaguchi, and Rachel Greenstadt
    USENIX Security Symposium (USENIX Security), 2015. (Acceptance rate: 15.7%)

15. Privacy Detective: Detecting Private Information and Collective Privacy Behavior in a Large Social Network
    Aylin Caliskan-Islam, Jonathan Walsh, and Rachel Greenstadt
    Workshop on Privacy in the Electronic Society (WPES), 2014. (Acceptance rate: 19%)

16. Doppelgänger Finder: Taking Stylometry To The Underground
    Sadia Afroz, **Aylin Caliskan-Islam**, Ariel Stolerman, Rachel Greenstadt, and Damon McCoy
    IEEE Symposium on Security and Privacy ("Oakland"), 2014. (Acceptance rate: 13.6%)
    Inaugural College of Computing & Informatics Day PhD Poster Award 2014

17. Approaches to Adversarial Drift
    Alex Kantchelian, Sadia Afroz, Ling Huang, **Aylin Caliskan-Islam**, Brad Miller, Michael Carl Tschantz, Rachel Greenstadt, Anthony Joseph, and J.D. Tygar
    6th ACM Workshop on Artificial Intelligence and Security (AISec), 2013.

18. How Privacy Flaws Affect Consumer Perception
    Sadia Afroz, **Aylin Caliskan-Islam**, Jordan Santell, Aaron Chapin, Rachel Greenstadt
    3rd Workshop on Socio-Technical Aspects in Security and Trust (STAST), 2013.

19. From Language to Family and Back: Native Language and Language Family Identification from English Text
    Ariel Stolerman, **Aylin Caliskan-Islam** and Rachel Greenstadt
    Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop (NAACL-SRW), 2013.

20. Use Fewer Instances of the Letter "i": Toward Writing Style Anonymization
    Andrew McDonald, Sadia Afroz, **Aylin Caliskan**, Ariel Stolerman and Rachel Greenstadt
    The 12th Privacy Enhancing Technologies Symposium (PETS), 2012. (Acceptance rate: 22.2%)
    Andreas Pfitzmann PETS Best Student Paper Award 2012

21. Translate once, translate twice, translate thrice and attribute: Identifying authors and machine translation tools in translated text
    Aylin Caliskan and Rachel Greenstadt
    6th IEEE International Conference on Semantic Computing (ICSC), 2012.

OTHER PUBLICATIONS

22. Detecting and mitigating bias in natural language processing
    Aylin Caliskan
    The Brookings Institution 2021

23. Comments in response to the National Institute of Standards and Technology Request for Information on Developing a Federal AI Standards Engagement Plan [Docket Number: [190312229–9229–01]]
    David Broniatowski, **Aylin Caliskan**, Valerie Reyna, and Reva Schwartz
    National Institute of Standards and Technology (NIST) White Paper, Jun 2019.

24. How do we decide how much to reveal? (Hint: Our privacy behavior might be socially constructed.)
    Aylin Caliskan-Islam
    Special Issue on Security, Privacy, and Human Behavior, ACM Computers & Society, Feb 2015.

25. ENVOY: Exploration and Navigation Vehicle for geolOgY Arunkumar Byravan, **Aylin Caliskan**, Jonas Cleveland, Daniel Gilles, Jaimeen Kapadia, Theparit Peerasathien, Bharath Sankaran, Alex Tozzo
    University of Pennsylvania's Entry in NASA/NIA RASC-AL, 2011.

AWARDED GRANTS

- Grants from NSF, NIST, DARPA, Amazon, and Microsoft.

BLOG POSTS

- When coding style survives compilation: De-anonymizing programmers from executable binaries, Dec 29, 2015.
- We can de-anonymize programmers from coding style. What are the implications?, Feb 26, 2015.
- Anonymous programmers can be identified by analyzing coding style, Jan 21, 2015.
- How do we decide how much to reveal? (Hint: Our privacy behavior might be socially constructed.), Dec 24, 2014.

POSTERS

- Dennis Röllke, Aviel J Stein, Edwin Dauber, Mosfiqur Rahman, Michael J Weisman, Gregory G Shearer, Frederica Nelson, Aylin Caliskan, Richard Harang, Rachel Greenstadt. Stylometry of Author-Specific and Country-Specific Style Features in JavaScript. NDSS, 2018.
- Edwin Dauber, Aylin Caliskan, Richard Harang, Rachel Greenstadt. Git blame who?: Stylistic authorship attribution of small, incomplete source code fragments Authors. IEEE/ACM 40th International Conference on Software Engineering: Companion, 2018.
- Aylin Caliskan. Privacy and Security via Machine Learning and Natural Language Processing. Cybersecurity Retreat, Princeton University, Mar 31, 2016.
- Aylin Caliskan-Islam. Doppelgänger Finder: Taking Stylometry To The Underground. Computer Science PhD Open House, Drexel University, 2014.
  Inaugural CCI Day PhD Poster Award
- Aylin Caliskan Islam. Supervised Source Code Authorship Attribution. Grace Hopper Conference, 2013.
- Aylin Caliskan Islam. Supervised Source Code Authorship Attribution. IEEE Symposium on Security and Privacy, 2013.
- Sadia Afroz, Aylin Caliskan Islam, Jordan Santell, Aaron Chapin and Rachel Greenstadt. How Privacy Flaws Affect Consumer Perception. SCRUB-UC Berkeley Winter Retreat, 2013.
- Aylin Caliskan Islam. Are we there yet? Can we identify authors and machine translation tools in translated text? ACM Student Research Competition, 2012.
- Arunkumar Byravan, Aylin Caliskan, Jonas Cleveland, Daniel Gilles, Jaimeen Kapadia, Theparit Peerasathien, Bharath Sankaran, Alex Tozzo. ENVOY: Exploration and Navigation Vehicle for geolOgY. NASA/NIA RASC-AL Forum, 2011.

SELECTED INVITED TALKS

- Aylin Caliskan. Artificial Intelligence for Social Good: When Machines Learn Human-like Biases from Data
  Harvard University, University of Chicago, University of Washington, Oracle Labs. Spring 2021.
- Aylin Caliskan. Implications of Biased AI on Democracy, Equity, and Justice.
  Invited Keynote COLING Workshop on Natural Language Processing for Internet Freedom. Dec 12, 2020.
- Aylin Caliskan. Gender Equity. AI for Good Global Summit. 2020.
- Aylin Caliskan. Bias in AI and Digital Humanities. University of Pennsylvania. 2020.
- Aylin Caliskan. Bias in AI. NIST AI Workshop. August 18, 2020.
- Aylin Caliskan. Bias and AI Ethics. DefCon28 AI Village. August 9, 2020.
- Aylin Caliskan. Algorithmic Measures of Language Mirror Human Biases. Georgetown University. Jan 24, 2020.
- Aylin Caliskan. Algorithmic Mirrors of Human Biases. Virginia Tech. Nov 8, 2019.
- Aylin Caliskan. Algorithmic Measures of Language Mirror Human Biases.
  Society of Experimental Social Psychology (SESP). Oct 18, 2019.
  Accepted Symposium on Computer-Resident Language and Naturalistic Conversation as Windows Into Social Cognition with Mahzarin Banaji, Benedek Kurdi, and Thalia Wheatley.
- Aylin Caliskan. Human-like Bias in Machine Intelligence. George Washington University, SEH WOW Talk Series. Oct 2, 2019.
- Aylin Caliskan. AI for Social Good, Bias and Ethics Panel. WeCNLP Summit at Facebook. Sep 6, 2019.
- Aylin Caliskan. Algorithmic Mirrors of Human Biases. University of Chicago. Jun 3, 2019.
- Aylin Caliskan. Monitoring Hate Speech in the US Media. Workshop on Defining, Monitoring and Countering Hate Speech. George Washington University, School of Media and Public Affairs. May 29, 2019.
- Aylin Caliskan. Algorithmic Mirrors of Society. University of Maryland. Apr 11, 2019.

- Aylin Caliskan. Algorithmic Mirrors of Society. National Academy of Sciences. Feb 20, 2019.
- Aylin Caliskan. Bias in AI. Social Science Foo Camp at Facebook. Feb 4, 2019.
- Aylin Caliskan. Bias in Machine Learning. ACM & Women in Computer Science at GWU. Oct 12, 2018.
- Aylin Caliskan. AI & Equity. **Invited Lecture** MIT Media Lab. Sep 25, 2018.
- Aylin Caliskan. The Great Power of AI: Algorithmic Mirrors of Society. DEFCON. Aug 11, 2018.
- Aylin Caliskan & Rachel Greenstadt. De-anonymizing Programmers from Source Code and Binaries. DEFCON. Aug 10, 2018.
- Aylin Caliskan. The Great Power of AI: Algorithmic Mirrors of Individuals and Society
  Brown University, Duke University, ETH Zurich, George Washington University, Tufts University, University of Maryland, University of Virginia, and Yale University. Feb - Apr 2018.
- Aylin Caliskan. Beyond Big Data: What Can We Learn from AI Models?
  Invited Keynote AISec - CCS Workshop. Nov 3, 2017.
- Aylin Caliskan. Natural Language Processing and Privacy: A Double Edged Sword.
  Invited Keynote Infer - PETS Workshop. Jul 18, 2016.
- Aylin Caliskan-Islam. Code Stylometry and Programmer De-anonymization. Göttingen University, Dec 23, 2015.
- Aylin Caliskan-Islam. De-anonymizing Programmers via Code Stylometry. Cornell Systems Lunch, Dec 4, 2015.
- Aylin Caliskan-Islam. De-anonymizing Programmers via Code Stylometry
  Guest Lecture, ELE574-Communications Security and Privacy, Princeton University, Oct 20, 2015.
- Aylin Caliskan-Islam. Security Review of Digital Privacy and the Underground: Miscreant Activity in the Internet
  Guest Lecture, CS475-Computer and Network Security, Mar 6, 2014.

TALKS

- Aylin Caliskan. Algorithmic Measures of Language Mirror Human Biases and Widely Shared Associations. Santa Fe Institute's Working Group: "Language as a Window into Human Minds: Exploration with Computer-Resident Language," Apr 23, 2020.
- Aylin Caliskan. Tutorial on Distributional Semantics via Word Embeddings. Department of Psychology, Harvard University. Nov 9, 2019.
- Aylin Caliskan. Neural Networks for NLP. Natural Language Processing Lecture at George Washington University, Nov 5, 2019.
- Aylin Caliskan. A Story of Discrimination and Unfairness. 33C3 - Chaos Communication Congress, Dec 27, 2016.
- Aylin Caliskan-Islam. A Story of Discrimination and Unfairness: Implicit Bias Embedded in Language Models. HotPETS 2016 - PETS. Jul 22, 2016. **Best Talk Award**
- Aylin Caliskan-Islam. De-anonymizing Programmers and Code Stylometry - Large Scale Authorship Attribution from Source Code and Executable Binaries of Compiled Code. Princeton University CITP Luncheon Speaker Series, Mar 1, 2016.
- Aylin Caliskan-Islam. De-anonymizing Programmers. 32C3 - Chaos Communication Congress, Dec 29, 2015.
- Aylin Caliskan-Islam. Support Vector Machines, Kernel Methods, Random Forests, and Feature Projection. Guest Lecture, CS613-Machine Learning, Jan 28, 2015. **Nominated for the Teaching Assistant Excellence Award**
- Aylin Caliskan-Islam, Rachel Greenstadt, and Rebekah Overdorf. Source Code and Cross-Domain Stylometry. Role of Stylometry in Privacy. 31st Chaos Communication Congress, Dec 29, 2014.
- Sadia Afroz and Aylin Caliskan-Islam. Stylometry and Online Underground Markets. 29th Chaos Communication Congress, Dec 28, 2012.
- Aylin Caliskan-Islam. Quantifying the Translator Effect: Identifying authors and machine translation tools in translated text. Girl Geek Dinners Philly, Feb 18, 2011.

WORKSHOPS

- Organizing Committee, NSF Workshop: Promises and Pitfalls of Big Data Approaches to Intersectional Equity in STEM. 2021
- Co-Chair, Gender Breakthrough. AI for Good Global Summit. 2020
- Participant and Speaker, NSF Workshop: Fairness, Ethics, Accountability, and Transparency (FEAT). 2019
- Chair, Hands-on Tutorial: AI Fairness 360, ACM Conference on Fairness, Accountability, and Transparency (ACM FAT*). 2019

EDITORIAL SERVICE

- Founding Editorial Board Member, Springer Journal, AI and Ethics

GRANT REVIEWS AND PANELS

- NSF 2020 (two programs)
- NSF 2019
- NASA 2019

SERVICE

- arXiv Moderator for Computer Science - Computers and Society (cs.CY) 2017 - present
- Reviewer, The Journal of Moral Philosophy 2021
- Reviewer, Nature Machine Intelligence 2021
- Program Committee, USENIX Security Symposium (USENIX Security) 2021
- Reviewer, The Proceedings of the National Academy of Sciences (PNAS) 2020
- Program Committee, Conference on Empirical Methods in Natural Language Processing (EMNLP) 2020
- Program Committee, Thirty-fourth Conference on Neural Information Processing Systems (NeurIPS) 2020
- Program Committee, 34th AAAI Conference on Artificial Intelligence (AAAI) 2020
- Gender Breakthrough Co-Chair and Braintrust Member. XPRIZE & AI for Good Global Summit 2020
- Program Committee, IEEE Symposium on Security and Privacy (Oakland S&P) 2020
- Program Committee, Privacy Enhancing Technologies Symposium (PETS) 2020
- Program Committee, IEEE Big Data 2019
- Program Committee, ACM Conference on Fairness, Accountability, and Transparency (ACM FAT*) 2019
- Program Committee, IEEE Symposium on Security and Privacy (Oakland S&P) 2019
- Program Committee, Deep Learning and Security Workshop co-located with IEEE S&P (Oakland) 2019
- Program Committee, ACM Workshop on Artificial Intelligence and Security co-located with the 26th ACM CCS (AISEC) 2019
- Program Committee, IEEE European Symposium on Security and Privacy (Euro S&P) 2018
- Program Committee, 11th ACM Workshop on Artificial Intelligence and Security (AISec) 2018
- Program Committee, The ACM Conference on Computer and Communications Security (CCS) 2017
- Program Committee, Privacy Enhancing Technologies Symposium (PETS) 2015-2018
- Program Committee, Data and Algorithmic Transparency Workshop (DAT) 2016
- Program Committee, AAAI Fall Symposium on Privacy and Language Technologies (AAAI PLT) 2016
- Reviewer, IEEE Transactions on Information Forensics and Security (TIFS) 2016
- Reviewer, Information and Software Technology 2016
- Reviewer, IEEE Transactions on Information Forensics and Security (TIFS) 2015
- Program Committee, Workshop on Privacy in the Electronic Society (WPES) 2015, 2018
- Subreviewer, USENIX Security Symposium (USENIX Security) 2015
- Subreviewer, International Joint Conference on Artificial Intelligence (IJCAI) 2015
- Subreviewer, Workshop on Privacy in the Electronic Society (WPES) 2014
- Subreviewer, USENIX Security Symposium (USENIX Security) 2014
- Subreviewer, Privacy Enhancing Technologies Symposium (PETS) 2014
- Subreviewer, The Sixth ASE International Conference on Social Computing (SocialCom) 2014

Aylin Caliskan

TEACHING EXPERIENCE

- Machine Learning, Fall 2020, (5 Undergraduate and 61 Graduate Students) George Washington University
- Senior Design Team1, Web Plug-in to Detect Bias in Text: Chloe Hutchins, Monica Kavathekar, and Jennifer Wright. 2019-2020
- Senior Design Team2, Tool for Detecting Plagiarism in Source Code: Jenny Fisher, Chengkai Ji, and Jinwen Zhang. 2019-2020
- Bias in AI, Spring 2020. (6 Undergraduate and 16 Graduate Students) George Washington University
- Research Credits, Spring 2020. (3 Graduate Students) George Washington University
- Machine Learning, Fall 2019, (15 Undergraduate and 47 Graduate Students) George Washington University
- Research Credits, Fall 2019. (3 Graduate Students) George Washington University
- Dissertation Credits, Fall 2019. (1 Graduate Student) George Washington University
- Bias in AI, Spring 2019. (4 Undergraduate and 8 Graduate Students) George Washington University
- Research Credits, Spring 2019. (2 Graduate Students) George Washington University
- Dissertation Credits, Spring 2019. (1 Graduate Student) George Washington University
- Machine Learning, Fall 2018, (1 Undergraduate and 36 Graduate Students) George Washington University
- De-anonymization and Machine Learning's Role in Privacy and Security
  Lecturer, Building Trust in The Information Age, Italy, Summer 2016
  Lectured PhD students in privacy and security at the summer school.
- Computer and Network Security (CS475)
  Teaching Assistant, Drexel University, Spring 2015
  Advised students on course material and projects.
- Graduate-Level Machine Learning (CS613)
  Teaching Assistant, Drexel University, Winter 2015
  Taught primary lecture material and prepared assignments.
  Nominated for the Teaching Assistant Excellence Award
- Data Structures and Algorithms in Java (CIS121)
  Lead Teaching Assistant & Teaching Assistant, University of Pennsylvania, Fall 2009, Spring 2010, Fall 2010, Spring 2011
  Held weekly recitation, prepared exam and project materials.

ADVISING

- Robert Wolfe, PhD advisee at George Washington University and University of Washington, 2020 - present
- Shiva Omrani, Master's student at George Washington University, June 2020-present
- Ryan Steed, Undergraduate Honors student at George Washington University, October 2018-present
  George Washington University Sigelman Undergraduate Research Enhancement Award for his work on Bias in AI
  Now a PhD student at Carnegie Mellon University's Heinz College of Information Systems & Public Policy
- Aviel Stein, PhD student at Drexel University, September 2017-present
- Abby Van Soest, senior thesis at Princeton University, September 2016-May 2017
- Neamah S. Abdulhusein, senior thesis at Princeton University, September 2016-May 2017
- Andrew Liu, undergraduate research intern working on source code authorship attribution at the US Army Research Laboratory, June 2014-November 2014
- Jordan Dantas, undergraduate research intern working on private information detection, September 2013-March 2014
- Senior Design Team, 5 seniors working on a sanitization project, September 2013-June 2014

UNIVERSITY COMMITTEES

- Faculty Search Committee, Dept. of Computer Science, George Washington University, 2020.
- Faculty Search Committee, Dept. of Engineering Management & Systems Engineering, George Washington University, 2020.
- Faculty Search Committee, Dept. of Computer Science, George Washington University, 2019.

DISSERTATION COMMITTEES

- Benjamin Bowman, Department of Electrical and Computer Engineering, George Washington University, 2020.
- Hong Xuan, Department of Computer Science, George Washington University, 2020.
- Pedram Hosseini, Department of Computer Science, George Washington University, 2020.
- Shabnam Tafreshi, Department of Computer Science, George Washington University, 2020.
- Xiao Xiao, Department of Computer Science, George Washington University, 2020.
- Siyuan Huang, Department of Computer Science, George Washington University, 2020. *(Master's Thesis)*
- Nada Almarwani, Department of Computer Science, George Washington University, 2019.
- Sawsan Alqahtani, Department of Computer Science, George Washington University, 2019.
- Fahad Ghamdi, Department of Computer Science, George Washington University, 2019.
- Yao Lu, Department of Computer Science, George Washington University, 2019.

ACTIVITIES

- Faculty Advisor, Women in Computer Science, George Washington University (2019 - 2021)
- President, Drexel Chapter of Upsilon Pi Epsilon (2012-2013)
- Vice President, Drexel University Computer Science Graduate Student Council (2011-2012)
- Member, Drexel University Women in Computing Society (WiCS) (2011-2012)
- Founding Member and Teleoperation, UPenn NASA Rasc-Al Robo Ops Team (2011)
- Treasurer and Vision Optimization, UPennalizers RoboCup Team (2009-2011)

PRESS COVERAGE

- Based on exact title search: "Semantics derived automatically from language corpora contain human-like biases." by Altmetric. Summary - December 2020: **Picked up by 97 news outlets**, Blogged by 32, 5 Policy Sources, Tweeted by 1,569, On 10 Facebook pages, 2 Wikipedia pages, Mentioned in 8 Google+ posts, Reddited by 2, Highlighted by 1 platform, On 1 videos, 1,016 readers on Mendeley, 8 readers on CiteULike. **Attention Score: 99th percentile, #29 of 267,043 outputs of similar age.**
- Uber Faces Civil Rights Lawsuit Alleging 'Racially Biased' Driver Ratings. Rachel Sandler, Forbes, 10.26.2020
- Was your Uber, Lyft fare high because of algorithm bias? Coral Murphy, USA TODAY, 07.22.2020.
- Researchers find racial discrimination in 'dynamic pricing' algorithms used by Uber, Lyft, and others. Kyle Wiggers. VentureBeat, 6.12.2020
- AI analyzed 3.3 million scientific abstracts and discovered possible new materials by Karen Hao. MIT Technology Review, 07.09.2019.
- AI Voice Assistants Reinforce Gender Biases, U.N. Report Says by Mahita Gajanan. Time, 05.22.2019.
- Why It's Dangerous For AI To Regulate Itself by Kayvan Alikhani. Forbes, 05.22.2019.
- Developing a moral compass from human texts by Patrick Bal. Technische Universitat Darmstadt, 02.07.2019.
- Yes, artificial intelligence can be racist by Brian Resnic. Vox, 01.24.2019.
- Mental health and artificial intelligence: losing your voice While we still can, let us ask, "Will AI exacerbate discrimination?" by Dan McQuillan. Open Democracy, 11.12.2018.
- Alexa and Google Home are no threat to regional accents – here's why by Erin Carrie. The Conversation, 08.21.2018.
- Machine Learning Can Identify the Authors of Anonymous Code: Even Anonymous Coders Leave Fingerprints by Louise Matsakis. Wired, 08.10.2018.
- AI Without Borders: How To Create Universally Moral Machines by Abinash Tripathy. Forbes, 04.11.2018.
- Princeton researchers discover why AI become racist and sexist by Annalee Newitz. Ars Technica, 04.18.2017.
- Training AI robots to act 'human' makes them sexist and racist by Mike Wehner. New York Post, 04.17.2017.
- How artificial intelligence learns how to be racist Simple: It's mimicking us. by Brian Resnick. Vox, 04.17.2017.
- When Artificial Intelligence = Not Enough Intelligence. by Michael Eric Ross. Omni, 04.16.2017.

- L'intelligence artificielle reproduit aussi le sexisme et le racisme des humains. by Morgane Tual. Le Monde, 04.15.2017.

- What would make a computer biased? Learning a language spoken by humans. by Melissa Healy. Los Angeles Times, 04.14.2017.

- A.I. Is Just as Sexist and Racist as We Are, and It's All Our Fault. by Peter Hess. Inverse, 04.14.2017.

- Podcast: Watching shoes untie, Cassini's last dive through the breath of a cryovolcano, and how human bias influences machine learning. by Sarah Crespi, David Grimm. Science Magazine, 04.13.2017.

- Even artificial intelligence can acquire biases against race and gender. by Matthew Hutson. Science Magazine, 04.13.2017.

- AI programs exhibit racial and gender biases, research reveals. by Hannah Devlin. The Guardian, 04.13.2017.

- AI Learns Gender and Racial Biases from Language. by Jeremy Hsu. IEEE Spectrum, 04.13.2017.

- Computers, Artificial Intelligence Show Bias and Prejudice, Too. by Maggie Fox. NBC News, 04.13.2017.

- AI robots learning racism, sexism and other prejudices from humans, study finds. by Ian Johnston. The Independent, 04.13.2017.

- Robots are learning to be racist AND sexist: Scientists reveal how AI programs exhibit human-like biases. by Stacy Liberatore. Daily Mail, 04.13.2017.

- Just like humans, artificial intelligence can be sexist and racist. by Matthew Burges. Wired UK, 04.13.2017.

- Scientists Taught A Robot Language. It Immediately Turned Racist. by Nidhi Subbaraman. BuzzFeed, 04.13.2017.

- AI picks up racial and gender biases when learning from what humans write. by Angela Chen. The Verge, 04.13.2017.

- Bad News: Artificial Intelligence Is Racist, Too. by Stephanie Pappas. Live Science, 04.13.2017.

- Surprise! AI can learn gender and race stereotypes, just like us. by Rebecca Ruiz. Mashable, 04.13.2017.

- Was passiert, wenn KIs unsere schlechten Eigenschaften übernehmen? by Anna Schughart. Wired, 02.24.2017.

- How to Fix Silicon Valley's Sexist Algorithms by Will Knight. MIT Technology Review, 11.23.2016.

- Bias in the machine: Internet algorithms reinforce harmful stereotypes by Bennett McIntosh. Discovery: Research at Princeton, 11.15.2016.

- Bath Researcher Shows Machines Can Be Prejudiced Too by Nick Flaherty. ACM TechNews, 10.21.2016.

- Artificial Intelligence Will Be as Biased and Prejudiced as Its Human Creators by Nathan Collins. Pacific Standard, 09.01.2016.

- It's Our Fault That AI Thinks White Names Are More 'Pleasant' Than Black Names by Jordan Pearson. Motherboard, 08.26.2016.

- CSI: Cyber-Attack Scene Investigation–a Malware Whodunit by Larry Greenemeier. Scientific American, 01.28.2016.

- Malicious coders will lose anonymity as identity-finding research matures by Joyce P. Brayboy. U.S. Army Research Laboratory and Comunications of the ACM 01.15.2016.

- 32C3: Nicht nur der Quellcode verrät seine Schöpfer by Kristian Kißling. Linux Magazin, 01.06.2016.

- De-Anonymizing Users from their Coding Styles by Bruce Schneier. Schneier on Security, 01.04.2016.

- Forget anonymity, we can remember you wholesale with machine intel, hackers warned by Alexander J. Martin. The Register, 12.31.2015.

- This Drexel researcher can identify you based on how you write code by Juliana Reyes. Technically Philly, 03.24.2015.

- Drexel researchers track 'cyber fingerprints' to identify programmers by Jessica McDonald. Newsworks, and two versions of interviews aired on WHYY Radio (NPR), 03.09.2015.

- Doctoral candidate uses coding style to identify hackers by Eric Birkhead. The Triangle, 02.28.2015.

- Dusting for Cyber Fingerprints: Computer Scientists Use Coding Style to Identify Anonymous Programmers by Britt Faulstick. Drexel Now, 02.26.2015.

- CSI Computer Science: Your coding style can give you away by Phil Johnson. ITWorld, 01.28.2015

- Your anonymous code contributions probably aren't: boffins by Richard Chirgwin. The Register, 01.22.2015.

- Open Campus initiative brings natural language processing to cyber research by United States Army Research Laboratory, 09.30.2014.

- Stylometric analysis to track anonymous users in the underground by Pierluigi Paganini. Security Affairs blog, 01.10.2013.

- Linguistics identifies anonymous users by Darren Pauli. SC Magazine (Australian edition), 01.09.2013.

- Students release stylometry tools by Helen Nowotnik. The Triangle, 01.13.2012.

- Software Helps Identify Anonymous Writers or Helps Them Stay That Way by Nicole Perlroth. New York Times Bits Blog. 01.03.2012. (An overview of Anonymouth and JStylo tool releases)

- Wer Hemingway imitiert, schreibt anonym by Philipp Elsbrock. Der Spiegel. 12.30.2011.
- State of Adversarial Stylometry: can you change your prose-style? by Cory Doctorow. Boing Boing. 12.29.2011. (An overview of our talk at 28C3, Anonymouth, and using it in fiction writing)
- Authorship recognition software from Drexel University lab to be released December by Christopher Wink. Technically Philly, 11.15.2011. (Announcing our upcoming release of JStylo & Anonymouth)
- RoboCup teams worldwide strive to bend it like C3PO by Tom Avril. The Philadelphia Inquirer, 06.10.2010.