

Fernando Diaz

diazf@acm.org
http://84r.io

March 12, 2026

Summary

I work on the design and evaluation of large-scale information access and AI systems. Across senior roles in industry and as a faculty member, my work has established algorithmic and evaluation frameworks and research directions now used in search, recommendation, and retrieval-augmented systems, influencing both academic discourse and production practice. I have led collaborations between research and engineering teams, built and guided research organizations, and contributed to community efforts shaping how AI systems are evaluated. My work integrates quantitative, qualitative, and mixed methods to improve system robustness.

Positions

Carnegie Mellon University
Pittsburgh, PA

Associate Professor with tenure
September 2023-present

I lead a research program on the evaluation and design of large-scale AI and information access systems, spanning foundational work in AI evaluation and the cultural and societal behavior of AI. My work has contributed to widely used methodologies for assessing retrieval-augmented and generative systems, influencing both research practice and production deployment. I collaborate across disciplines and with industry, and I lead and co-lead international projects on AI and culture. I train PhD students for roles in academia and industry and developed the first graduate course on quantitative evaluation of language technologies.

Google Research
Montréal, QC
Pittsburgh, PA

Staff Research Scientist
September 2020-March 2023
April 2023-present

I work on the evaluation and design of search and recommender systems, with a focus on responsible AI in production environments. My work contributed to methodological and empirical advances in how large-scale systems are evaluated, and included close collaboration with colleagues in HCI and related areas on the human and societal dimensions of AI systems.

In parallel, I developed a multidisciplinary research program examining the longer-term cultural and societal effects of music creation, recommendation, and consumption.

Microsoft Research
Montréal, QC

Senior Principal Research Manager, Assistant Research Director
June 2018-July 2020

I built and led a multidisciplinary research group focused on fairness, accountability, transparency, and ethics in AI systems, establishing a research agenda that informed both research and product at Microsoft. The group combined foundational research with close collaboration with engineering teams, contributing to evaluation and ranking approaches used in production systems such as Bing.

In parallel, I maintained a research program in core information retrieval and initiated new

work on the broader cultural and societal impacts of recommendation and creative AI systems. As Assistant Research Director, I helped oversee a portfolio of approximately twenty-five researchers across NLP, machine learning, and responsible AI, contributing to hiring, strategy, and organizational direction. Work during this time received multiple awards for foundational contributions to responsible AI.

Spotify
New York, NY

Director of Research
March 2017-May 2018

I built and led a research organization focused on search and recommender systems, including metrics and evaluation. I scaled the team from three researchers in a single location to approximately twenty-five across three locations, establishing research leadership roles and lab directors to support growth. I helped define how research informed product development, shaping evaluation and experimentation practices and longer-term technical direction within the search and recommendations organization.

As a member of the search and personalization leadership team, I contributed to strategic planning and the technical roadmap for recommender systems, helping guide product and infrastructure investments over a multi-year horizon. I also developed collaborations with the academic community and strengthened the role of research within Spotify's engineering and product culture.

Microsoft Research
New York, NY

Senior Researcher
May 2012-March 2017

I conducted basic research and co-founded early work on fairness, accountability, transparency, and ethics (FATE) in AI, helping establish a research direction that later grew into a broader organizational effort. My work spanned information retrieval, large-scale attention and interaction modeling, fairness in machine learning, text summarization, and deep learning for search, combining foundational research with collaboration on production-scale systems.

I worked closely with product and engineering teams to translate research into system design and evaluation practices, and collaborated extensively with academic partners. Work during this time received multiple awards for foundational contributions to crisis informatics.

Yahoo! Research
New York, NY

Senior Research Scientist
January 2008-March 2012

I led and contributed to research in federated search, time-sensitive web search, large-scale user attention and interaction modeling, match-making systems, and multi-document summarization. This work combined foundational advances in information retrieval with the development of production-quality systems and close collaboration with applied research and engineering teams.

I helped transfer research into deployed search technologies and evaluation practices, and engaged actively with the broader research community through publications and collaboration. This research was recognized at multiple conferences for contributions to information retrieval.

Institute for Pure and Applied Mathematics
Los Angeles, CA

Visiting Fellow
September 2007-December 2007

I was a visiting fellow at UCLA for the program in "Mathematics of Knowledge and Search Engines".

University of Massachusetts
Amherst, MA

Research Assistant
September 2001-August 2007

I conducted dissertation research focused on machine learning approaches to information retrieval under the supervision of James Allan.

Overture Research
Pasadena, CA

Intern
June 2003-December 2003

I collaborated with Dr. Rosie Jones to develop tools for the analysis and classification of search queries into temporal categories.

Sony Computer Science Laboratories
Tokyo, Japan

Intern
June 2001-August 2001

As part of participation in the NSF Summer Institute in Japan, I worked with Dr. Hitoshi Iida and Dr. Koiti Hasida of Sony Computer Science Laboratories on the construction of a dialog-based information retrieval system.

University of Massachusetts
Amherst, MA

Research Assistant
September 2000-May 2001

I collaborated with Professors Victor Lesser and Beverly Woolf on a project which leveraged information retrieval techniques in a distributed educational system. We developed a system to infer symbolic and planning primitives from statistical text analysis.

Bell & Howell Information and Learning
Ann Arbor, MI

Research Assistant
May 1999-August 2000

I worked on implementing information retrieval methods in a production search system indexing both native-text and scanned document collections. My work included system design, engineering, and statistical evaluation of retrieval effectiveness.

**Academic
Affiliation**

Carnegie Mellon University
Pittsburgh, PA

Associate Professor with tenure
August 2023-present

McGill University
Montréal, QC

Adjunct Professor
January 2019-present

NYU Courant Institute of Mathematical Sciences
New York, NY

Adjunct Professor
January 2013-December 2016

NYU Tandon School of Engineering
New York, NY

Adjunct Professor
January 2011-May 2011

Publications

Metrics

	<i>ACM DL</i>	<i>Scopus</i>	<i>Google Scholar</i>	<i>Semantic Scholar</i>
articles	105	121	199	141
citations	3842	6646	13358	9258
citations/article	36.6	54.9	67.1	65.7
h-index	33	42	56	46

Conference Papers

Abinay Reddy Naini, **Fernando Diaz**, and Carlos Busso. RankList – A Listwise Preference Learning Framework for Predicting Subjective Preferences. In *Proceedings of the 40th Annual AAAI Conference on Artificial Intelligence*, 2026

Rikiya Takehi, **Fernando Diaz**, and Tetsuya Sakai. Diversification as Risk Minimization. In *Proceedings of the Nineteenth ACM International Conference on Web Search and Data Mining*, 2026. **Best Paper**

Alexandra Olteanu, Su Lin Blodgett, Agathe Balayn, Angelina Wang, **Fernando Diaz**, Flavio du Pin Calmon, Margaret Mitchell, Michael Ekstrand, Reuben Binns, and Solon Barocas. Rigor in AI: Doing Rigorous AI Work Requires a Broader, Responsible AI-Informed Conception of Rigor. In *Advances in Neural Information Processing Systems*, 2025

Jushaan Singh Kalra, Xinran Zhao, To Eun Kim, Fengyu Cai, **Fernando Diaz**, and Tongshuang Wu. MoR: Better Handling Diverse Queries with a Mixture of Sparse, Dense, and Human Retrievers. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 2025

To Eun Kim and **Fernando Diaz**. Towards Fair RAG: On the Impact of Fair Ranking in Retrieval-Augmented Generation. In *Proceedings of the 2025 International ACM SIGIR Conference on Innovative Concepts and Theories in Information Retrieval (ICTIR)*, 2025

Yifan He, To Eun Kim, **Fernando Diaz**, Jaime Arguello, and Bhaskar Mitra. Tip of the Tongue Query Elicitation for Simulated Evaluation. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2025

Athiya Deviyani and **Fernando Diaz**. Contextual Metric Meta-Evaluation by Measuring Local Metric Accuracy. In *Findings of the Association for Computational Linguistics: NAACL 2025*, 2025

Shaily Bhatt and **Fernando Diaz**. Extrinsic Evaluation of Cultural Competence in Large Language Models. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen, editors, *Findings of the Association for Computational Linguistics: EMNLP 2024*, November 2024

Fernando Diaz. Pessimistic Evaluation. In *Proceedings of the Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, 2024

Negar Arabzadeh, **Fernando Diaz**, and Junfeng He. Offline Evaluation of Set-Based Text-to-Image Generation. In *Proceedings of the Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, 2024

Haolun Wu, Ofer Meshi, Masrour Zoghi, **Fernando Diaz**, Xue Liu, Craig Boutilier, and Maryam Karimzadehgan. Density-based User Representation using Gaussian Process Regression for Multi-interest Personalized Retrieval. In *Advances in Neural Information Processing Systems*, 2024

- Fernando Diaz and Michael Madaio. Scaling Laws Do Not Scale. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7, 2024
- Ali Vardasbi, Maarten de Rijke, Fernando Diaz, and Mostafa Dehghani. The Impact of Group Membership Bias on the Quality and Fairness of Exposure in Ranking. In *Proceedings of the 47th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2024
- Lauren Wilcox, Robin Brewer, and Fernando Diaz. AI Consent Futures: A Case Study on Voice Data Collection with Clinicians. *Proc. ACM Hum.-Comput. Interact.*, 7(CSCW2), oct 2023.
Honorable Mention
- Andres Ferraro, Gustavo Ferreira, Fernando Diaz, and Georgina Born. Measuring Commonality in Recommendation of Cultural Content: Recommender Systems to Enhance Cultural Citizenship. In *Proceedings of the 16th ACM Conference on Recommender Systems*, 2022
- Fernando Diaz and Andres Ferraro. Offline Retrieval Evaluation Without Evaluation Metrics. In *Proceedings of the 45th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2022
- Hamed Zamani, Fernando Diaz, Mostafa Dehghani, Donald Metzler, and Michael Bendersky. Retrieval-Enhanced Machine Learning. In *Proceedings of the 45th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2022
- Haolun Wu, Bhaskar Mitra, Chen Ma, Fernando Diaz, and Xue Liu. Joint Multisided Exposure Fairness for Recommendation. In *Proceedings of the 45th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2022
- Filip Radlinski, Krisztian Balog, Fernando Diaz, Lucas Dixon, and Ben Wedin. On Natural Language User Profiles for Transparent and Scrutable Recommendation. In *Proceedings of the 45th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2022
- Divya Shanmugam, Fernando Diaz, Samira Shabaniyan, Michele Finck, and Asia Biega. Learning to Limit Data Collection via Scaling Laws: A Computational Interpretation for the Legal Principle of Data Minimization. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, 2022
- Ruohan Li, Jianxiang Li, Bhaskar Mitra, Fernando Diaz, and Asia J. Biega. Exposing Query Identification for Search Transparency. In *Proceedings of the 31st International Conference on World Wide Web*, 2022
- Ramya Srinivasan, Emily Denton, Jordan Famularo, Negar Rostamzadeh, Fernando Diaz, and Beth Coleman. Artsheets for Art Datasets. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, 2021
- Omer Kirnap, Fernando Diaz, Asia Biega, Michael Ekstrand, Ben Carterette, and Emine Yilmaz. Estimation of Fair Ranking Metrics with Incomplete Judgments. In *Proceedings of the 30th International Conference on World Wide Web*, 2021
- Jaime Arguello, Adam Ferguson, Emery Fine, Bhaskar Mitra, Hamed Zamani, and Fernando Diaz. Tip of the Tongue Known-Item Retrieval: A Case Study in Movie Identification. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval*, 2021
- Ronald Robertson, Alexandra Olteanu, Fernando Diaz, Milad Shokouhi, Peter Bailey, and Shashank Jain. Characterizing Problematic Email Reply Suggestions. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021

- Fernando Diaz**, Bhaskar Mitra, Michael D. Ekstrand, Asia J. Biega, and Ben Carterette. Evaluating Stochastic Rankings with Expected Exposure. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, 2020. **Best Paper Nomination**
- Alexandra Olteanu, **Fernando Diaz**, and Gabriella Kazai. When Are Search Completion Suggestions Problematic? *Proc. ACM Hum.-Comput. Interact.*, 4(CSCW2), October 2020. **Honorable Mention**
- Hamed Zamani, Bhaskar Mitra, Everest Chen, Gord Lueck, **Fernando Diaz**, Paul Bennett, Nick Craswell, and Susan Dumais. Analyzing and Learning from User Interactions for Search Clarification. In *Proceedings of the 43rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2020
- Asia J. Biega, Peter Potash, Hal Daumé III, **Fernando Diaz**, and Michèle Finck. Operationalizing the Legal Principle of Data Minimization for Personalization. In *Proceedings of the 43rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2020
- Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and **Fernando Diaz**. Towards a Fair Marketplace: Counterfactual Evaluation of the trade-off between Relevance, Fairness and Satisfaction in Recommendation Systems. In *Proceedings of the 27th ACM conference on Information and knowledge management*, 2018
- Jean Garcia-Gathright, Brian St. Thomas, Christine Hosey, Zahra Nazari, and **Fernando Diaz**. Understanding and Evaluating User Satisfaction with Music Discovery. In *Proceedings of the 41st International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2018
- Rishabh Mehrotra, Amit Sharma, Ashton Anderson, **Fernando Diaz**, Hanna Wallach, and Emine Yilmaz. Auditing Search Engines for Differential Satisfaction Across Demographics. In *Proceedings of the 26th International Conference on World Wide Web*, 2017
- Omar Alonso, Serge-Eric Tremblay, and **Fernando Diaz**. Automatic Generation of Event Timelines from Social Data. In *Proceedings of the 2017 ACM on Web Science Conference*, 2017
- Bhaskar Mitra, **Fernando Diaz**, and Nick Craswell. Learning to Match Using Local and Distributed Representations of Text for Web Search. In *Proceedings of the 26th International Conference on World Wide Web*, 2017
- Matthew Ekstrand-Abueg, Richard McCreddie, Virgil Pavlu, and **Fernando Diaz**. A Study of Realtime Summarization Metrics. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, 2016
- Rahul Goel, Sandeep Soni, Naman Goyal, John Paparrizos, Hanna Wallach, **Fernando Diaz**, and Jacob Eisenstein. *The Social Dynamics of Language Change in Online Networks*. 2016
- Fernando Diaz**. Learning to Rank with Labeled Features. In *Proceedings of the 2016 ACM International Conference on the Theory of Information Retrieval*, 2016
- Fernando Diaz**, Bhaskar Mitra, and Nick Craswell. Query Expansion with Locally-Trained Word Embeddings. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, 2016
- Fernando Diaz**, Qi Guo, and Ryen W. White. Search Result Prefetching Using Cursor Movement. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2016

- Christopher Kedzie, **Fernando Diaz**, and Kathleen McKeown. Real-Time Web Scale Event Summarization Using Sequential Decision Making. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, 2016
- Fernando Diaz**. Pseudo-Query Reformulation. In *Proceedings of the 38th European Conference on IR Research*, 2016
- Jaime Arguello, Sandeep Avula, and **Fernando Diaz**. Using Query Performance Predictors to Improve Spoken Queries. In *Proceedings of the 38th European Conference on IR Research*, 2016
- Fernando Diaz**. Condensed List Relevance Models. In *Proceedings of the 2015 International Conference on The Theory of Information Retrieval*, May 2015
- Christopher Kedzie, Kathleen McKeown, and **Fernando Diaz**. Predicting Salient Updates for Disaster Summarization. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, July 2015
- Pavel Metrikov, **Fernando Diaz**, Sébastien Lahaie, and Justin Rao. Whole Page Optimization: How Page Elements Interact with the Position Auction. In *EC '14: Proceedings of the fifteenth ACM conference on Electronic commerce*, 2014
- Alexandra Olteanu, Carlos Castillo, **Fernando Diaz**, and Sarah Vieweg. CrisisLex: A Lexicon for Collecting and Filtering Microblogged Communications in Crises. In *Proceedings of The 8th International AAI Conference on Weblogs and Social Media*, 2014
- Milad Shokouhi, Rosie Jones, Umut Ozertem, Karthik Raghunathan, and **Fernando Diaz**. Mobile Query Reformulations. In *Proceedings of the 37th Annual ACM SIGIR Conference*, 2014
- Peter B. Golbus, Imed Zitouni, Jin Young Kim, Ahmed Hassan, and **Fernando Diaz**. Contextual and Dimensional Relevance Judgments for Reusable SERP-level Evaluation. In *Proceedings of the 23rd International Conference on World Wide Web*, 2014
- Fernando Diaz**, Ryen W. White, Georg Buscher, and Dan Liebling. Robust Models of Mouse Movement on Dynamic Web Search Results Pages. In *Proceedings of the 22nd ACM conference on Information and knowledge management (CIKM 2013)*, 2013
- Muhammad Imran, Shady Elbassuoni, Carlos Castillo, **Fernando Diaz**, and Patrick Meier. Extracting Information Nuggets from Disaster-Related Messages in Social Media. In *10th International Conference on Information Systems for Crisis Response and Management*, 2013. **Best Paper**
- Qi Guo, **Fernando Diaz**, and Elad Yom-Tov. Updating Users about Time Critical Events. In Pavel Serdyukov, Pavel Braslavski, Sergei O. Kuznetsov, Jaap Kamps, Stefan Rüger, Eugene Agichtein, Ilya Segalovich, and Emine Yilmaz, editors, *Advances in Information Retrieval*, volume 7814. 2013
- Jaime Arguello, **Fernando Diaz**, and Jamie Callan. Learning to aggregate vertical results into web search results. In *Proceedings of the 20th ACM international conference on Information and knowledge management*, 2011
- Elad Yom-Tov and **Fernando Diaz**. Out of sight, not out of mind: on the effect of social and physical detachment on information need. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*, 2011. **Best Paper Nomination**

- Elad Yom-Tov and **Fernando Diaz**. Location and timeliness of information sources during news events. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*, 2011
- Jangwon Seo, **Fernando Diaz**, Evgeniy Gabrilovich, Vanja Josifovski, and Bo Pang. Generalized Link Suggestions via Web Site Clustering. In *Proceedings of the 20th International Conference on World Wide Web*, 2011
- Jaime Arguello, **Fernando Diaz**, Jamie Callan, and Ben Carterette. A methodology for evaluating aggregated search results. In *Proceedings of the 33rd European conference on Advances in information retrieval*, 2011. **Best Student Paper**
- Jing Bai, **Fernando Diaz**, Yi Chang, Zhaohui Zheng, and Keke Chen. Cross-market model adaptation with pairwise preference data for web search ranking. In *Proceedings of the 23rd International Conference on Computational Linguistics: Posters*, 2010
- Fernando Diaz**, Donald Metzler, and Sihem Amer-Yahia. Relevance and Ranking in Online Dating Systems. In *SIGIR '10: Proceeding of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, 2010. **ICML 2011 Invited Cross-Conference Presentation**
- Jaime Arguello, **Fernando Diaz**, and Jean-François Paiement. Vertical selection in the presence of unlabeled verticals. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, 2010
- Anlei Dong, Ruiqiang Zhang, Pranam Kolari, Jing Bai, **Fernando Diaz**, Yi Chang, Zhaohui Zheng, and Hongyuan Zha. Time is of the essence: improving recency ranking using Twitter data. In *WWW '10: Proceedings of the 19th international conference on World wide web*, 2010
- Anlei Dong, Yi Chang, Zhaohui Zheng, Gilad Mishne, Jing Bai, Ruiqiang Zhang, Karolina Buchner, Ciya Liao, and **Fernando Diaz**. Towards recency ranking in web search. In *WSDM '10: Proceedings of the third ACM international conference on Web search and data mining*, 2010
- Jaime Arguello, Jamie Callan, and **Fernando Diaz**. Classification-Based Resource Selection. In *CIKM '09: Proceeding of the 18th ACM conference on Information and knowledge management*, 2009
- Ahmed Hassan, Rosie Jones, and **Fernando Diaz**. A case study of using geographic cues to predict query news intent. In *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 2009
- Jaime Arguello, **Fernando Diaz**, Jamie Callan, and Jean-François Crespo. Sources of Evidence for Vertical Selection. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, 2009. **Best Paper**
- Fernando Diaz** and Jaime Arguello. Adaptation of Offline Vertical Selection Predictions in the Presence of User Feedback. In *SIGIR 2009*, 2009
- Fernando Diaz**. Integration of News Content Into Web Results. In *Proceedings of the Second ACM International Conference on Web Search and Data Mining*, 2009. **Best Paper**
- Fernando Diaz**. Improving relevance feedback in language modeling with score regularization. In *SIGIR '08: Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, 2008
- Fernando Diaz**. A method for transferring retrieval scores between collections with non-overlapping vocabularies. In *SIGIR '08: Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, 2008

- Fernando Diaz.** Theoretical bounds on and empirical robustness of score regularization to different similarity measures. In *SIGIR '08: Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, 2008
- Fernando Diaz.** Performance prediction using spatial autocorrelation. In *SIGIR '07: Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, 2007
- Fernando Diaz** and Donald Metzler. Pseudo-Aligned Multilingual Corpora. In Manuela M. Veloso, editor, *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence*, 2007
- Fernando Diaz** and Donald Metzler. Improving the estimation of relevance models using large external corpora. In *SIGIR '06: Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, 2006
- Fernando Diaz.** Regularizing ad hoc retrieval scores. In *CIKM '05: Proceedings of the 14th ACM international conference on Information and knowledge management*, 2005
- Fernando Diaz** and Rosie Jones. Using temporal profiles of queries for precision prediction. In *SIGIR '04: Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*, 2004
- Diane Kelly, **Fernando Diaz**, Nicholas J. Belkin, and James Allan. A User-Centered Approach to Evaluating Topic Models. In *26th European Conference on Information Retrieval Research*, 2004
- Fernando Diaz.** Using Wearable Computers to Construct Semantic Representations of Physical Spaces. In *Proceedings of the Sixth International Symposium on Wearable Computers*, 2002

Journal Articles

- Asia Biega, Georgina Born, **Fernando Diaz**, Mary L. Gray, and Rida Qadri. Towards a Multidisciplinary Vision for Culturally Inclusive Generative AI (Dagstuhl Seminar 25022). *Dagstuhl Reports*, 15(1), 2025
- Fernando Diaz**, Michael D. Ekstrand, and Bhaskar Mitra. Recall, Robustness, and Lexicographic Evaluation. *ACM Trans. Recomm. Syst.*, 4(1), July 2025
- Andres Ferraro, Gustavo Ferreira, **Fernando Diaz**, and Georgina Born. Measuring Commonality in Recommendation of Cultural Content to Strengthen Cultural Citizenship. *ACM Trans. Recomm. Syst.*, 2(1), mar 2024
- Michael D. Ekstrand, Ben Carterette, and **Fernando Diaz**. Distributionally-Informed Recommender System Evaluation. *ACM Trans. Recomm. Syst.*, 2(1), mar 2024
- Haolun Wu, Chen Ma, Bhaskar Mitra, **Fernando Diaz**, and Xue Liu. A Multi-Objective Optimization Framework for Multi-Stakeholder Fairness-Aware Recommendation. *ACM Trans. Inf. Syst.*, 41(2), dec 2022
- Nancy Baym, Rachel Bergmann, Raj Bhargava, **Fernando Diaz**, Tarleton Gillespie, David Hesmondhalgh, Elena Maris, and Christopher Persaud. Making Sense of Metrics in the Music Industries. *International Journal of Communication*, 15(0), 2021
- Georgina Born, Jeremy Morris, **Fernando Diaz**, and Ashton Anderson. Artificial Intelligence, Music Recommendation, and the Curation of Culture. *Schwartz Reisman Institute for Technology and Society White Paper*, 2021

- J. Shane Culpepper, **Fernando Diaz**, and Mark D. Smucker. Research Frontiers in Information Retrieval: Report from the Third Strategic Workshop on Information Retrieval in Lorne (SWIRL 2018). *SIGIR Forum*, 52(1), August 2018
- Alexandra Olteanu, Carlos Castillo, **Fernando Diaz**, and Emre Kiciman. Social Data: Biases, Methodological Pitfalls, and Ethical Boundaries. *Frontiers in Big Data*, 2, 2017
- Ryen W. White, **Fernando Diaz**, and Qi Guo. Search Result Prefetching on Desktop and Mobile. *ACM Trans. Inf. Syst.*, 35(3), May 2017
- Fernando Diaz**. Worst Practices for Designing Production Information Access Systems. *SIGIR Forum*, 50(1), June 2016
- Fernando Diaz**, Michael Gamon, Jake Hofman, Emre Kiciman, and David Rothschild. Online and Social Media Data As an Imperfect Continuous Panel Survey. *PLoS ONE*, 11(1), 2016
- Muhammad Imran, Carlos Castillo, **Fernando Diaz**, and Sarah Vieweg. Processing Social Media Messages in Mass Emergency: A Survey. *ACM Comput. Surv.*, 47(4), July 2015
- Fernando Diaz**. Experimentation Standards for Crisis Informatics. *SIGIR Forum*, 48(2), December 2014
- Daniel G. Goldstein, Siddharth Suri, R. Preston McAfee, Matthew Ekstrand-Abueg, and **Fernando Diaz**. The Economic and Cognitive Costs of Annoying Display Advertisements. *Journal of Marketing Research*, 51(6), December 2014. **Finalist: Paul E. Green Award**
- Hemant Purohit, Carlos Castillo, **Fernando Diaz**, Amit Sheth, and Patrick Meier. Emergency-Relief Coordination on Social Media: Automatically Matching Resource Requests and Offers. *First Monday*, 19(1), 2014
- Elad Yom-Tov and **Fernando Diaz**. The Effect of Social and Physical Detachment on Information Need. *ACM Trans. Inf. Syst.*, 31(1), January 2013
- Yi Chang, Anlei Dong, Pranam Kolari, Ruiqiang Zhang, Yoshiyuki Inagaki, **Fernando Diaz**, Hongyuan Zha, and Yan Liu. Improving Recency Ranking Using Twitter Data. *ACM Trans. Intell. Syst. Technol.*, 4(1), February 2013
- Fernando Diaz**. Regularizing query-based retrieval scores. *Information Retrieval*, 10(6), December 2007
- Rosie Jones and **Fernando Diaz**. Temporal profiles of queries. *ACM Trans. Inf. Syst.*, 25(3), July 2007

Chapters

- Jaime Arguello and **Fernando Diaz**. *Relevance Ranking of Vertical Search Engines*, chapter Vertical Selection and Aggregation. 2013

Books

- Michael D. Ekstrand, Anubrata Das, Robin Burke, and **Fernando Diaz**. *Fairness and Discrimination in Information Access Systems*. 2022

Thesis

Fernando Diaz. *Autocorrelation and Regularization of Query-Based Retrieval Scores.* PhD thesis, University of Massachusetts Amherst, February 2008

Workshop Papers

Jaime Arguello, **Fernando Diaz**, Maik Fröbe, To Eun Kim, and Bhaskar Mitra. Overview of the TREC 2025 Tip-of-the-Tongue Track. In *Proceedings of the Thirty-Fourth Text REtrieval Conference*, 2026

Jaime Arguello, Samarth Bhargav, **Fernando Diaz**, Evangelos Kanoulas, To Eun Kim, Yifan He, and Bhaskar Mitra. Overview of the TREC 2024 Tip-of-the-Tongue Track. In *Proceedings of the Thirty-Third Text REtrieval Conference*, 2025

Jaime Arguello, Samarth Bhargav, **Fernando Diaz**, Evangelos Kanoulas, and Bhaskar Mitra. Overview of the TREC 2023 Tip-of-the-Tongue Track. In *Proceedings of the Thirty-Second Text REtrieval Conference*, 2024

Fernando Diaz. Best-Case Retrieval Evaluation: Improving the Sensitivity of Reciprocal Rank with Lexicographic Precision. In *Proceedings of the 10th International Workshop on Evaluating Information Access co-located with the 17th NTCIR Conference on the Evaluation of Information Access Technologies (NTCIR 2023)*, 2023

Esther Rolf, Ben Packer, Alex Beutel, and **Fernando Diaz.** Striving for data-model efficiency: Identifying data externalities on group performance. In *Workshop on Trustworthy and Socially Responsible Machine Learning, NeurIPS 2022*, 2022

Asia J. Biega, **Fernando Diaz**, Michael D. Ekstrand, Sergey Feldman, and Sebastian Kohlmeier. Overview of the TREC 2020 Fair Ranking Track. In *The Twenty-Eighth Text REtrieval Conference (TREC 2020) Proceedings*, 2020

Asia J. Biega, **Fernando Diaz**, Michael D. Ekstrand, and Sebastian Kohlmeier. Overview of the TREC 2019 Fair Ranking Track. In *The Twenty-Eighth Text REtrieval Conference (TREC 2019) Proceedings*, 2019

Sarah Bird, Solon Barocas, Kate Crawford, **Fernando Diaz**, and Hanna Wallach. Exploring or Exploiting? Social and Ethical Implications of Autonomous Experimentation in AI. In *Workshop on Fairness, Accountability, and Transparency in Machine Learning*, 2016

Jimmy Lin, Adam Roegiest, Luchen Tan, Richard McCreadie, Ellen Voorhees, and **Fernando Diaz.** Overview of the TREC 2016 Real-Time Summarization Track. In *The 26th Text Retrieval Conference Proceedings (TREC 2016)*, 2016. Special Publication

David Abel, Alekh Agarwal, **Fernando Diaz**, Akshay Krishnamurthy, and Robert E. Schapire. Exploratory Gradient Boosting for Reinforcement Learning in Complex Domains. In *ICML 2016 Workshop on Abstraction in Reinforcement Learning*, 2016

Javed Aslam, **Fernando Diaz**, Matthew Ekstrand-Abueg, Richard McCreadie, Virgi Pavlu, and Tetsuya Sakai. TREC 2015 Temporal Summarization Track Overview. In *The 24th Text Retrieval Conference Proceedings (TREC 2015)*, 2015. Special Publication

Fernando Diaz. Experimentation Standards for Crisis Informatics. In *KDD 2014 Workshop on Data Science for Social Good*, 2014

Javed Aslam, Matthew Ekstrand-Abueg, Virgi Pavlu, **Fernando Diaz**, Richard McCreadie, and Tetsuya Sakai. TREC 2014 Temporal Summarization Track Overview. In *The 23rd Text Retrieval Conference Proceedings (TREC 2014)*, 2014. Special Publication

Javed A. Aslam, **Fernando Diaz**, Matthew Ekstrand-Abueg, Virgil Pavlu, and Tetsuya Sakai. TREC 2013 Temporal Summarization. In *Proceedings of the Twenty-Second Text REtrieval Conference (TREC 2013)*, 2013

Muhammad Imran, Shady Elbassuoni, Carlos Castillo, **Fernando Diaz**, and Patrick Meier. Practical extraction of disaster-relevant information from social media. In *Proceedings of the 22nd international conference on World Wide Web companion*, 2013

Annie Louis, Eric Crestan, Youssef Billawala, Rao Shen, **Fernando Diaz**, and Jean-François Crespo. Use of Query Similarity for Improving Presentation of News Verticals. In *Proceedings of the First International Workshop on Searching and Integrating New Web Data Sources - Very Large Data Search, Seattle, WA, USA, September 2, 2011*, 2011

Rosie Jones, Ahmed Hassan, and **Fernando Diaz**. Geographic Features in Web Search Retrieval. In *Proceedings of the 2Nd International Workshop on Geographic Information Retrieval*, 2008

Donald Metzler, **Fernando Diaz**, Trevor Strohman, and W. B. Croft. UMass at Robust 2005: Using Mixtures of Relevance Models for Query Expansion. In *The Fourteenth Text REtrieval Conference (TREC 2005) Notebook*, 2005

Fernando Diaz and James Allan. When Less is More: Relevance Feedback Falls Short and Term Expansion Succeeds at HARD 2005 (Notebook Version). In *The Fourteenth Text REtrieval Conference (TREC 2005) Notebook*, 2005

Nasreen Abdul-Jaleel, James Allan, W. Bruce Croft, **Fernando Diaz**, Leah Larkey, Xiaoyan Li, Donald Metzler, Mark D. Smucker, Trevor Strohman, Howard Turtle, and Courtney Wade. UMass at TREC 2004: Novelty and HARD. In *Online Proceedings of 2004 Text REtrieval Conference*, 2004

Education

PhD	University of Massachusetts Amherst Computer Science "Regularizing Query-Based Retrieval Scores" James Allan, W. Bruce Croft, Sridhar Mahadevan, John Staudenmayer	2008
MS	University of Massachusetts Amherst Computer Science "Browsing-Based User Language Models" James Allan	2004
BS	University of Michigan Ann Arbor Computer Science	1998

Teaching

Academic

Quantitative Evaluation of Language Technologies

Carnegie Mellon University (Language Technologies Institute) 2024

Graduate level course on measurement and experimentation of language technologies in offline and online environments. Prepared all material, lectures, and homeworks.

Web Search Engines

New York University (Courant) 2013-2016

Co-instructed (with Cong Yu) a graduate level course on web search engines. Classes consisted of lectures with evaluation based on homeworks, exams, and a final project. I presented half of the lectures. Class sizes ranged from thirty to fifty.

Web Search

Asian Summer School in Information Access

Summer 2013

Graduate level lecture on web search engines. Class size was 63 students.

Experimental Design for Information Systems

University of Trento

Summer 2012

Co-instructed (with Cong Yu) a graduate level course on information retrieval and data mining. Classes included a combination of lectures and paper discussion. I presented half of the lectures. Class size was 20 students.

Web Search Engines

New York University (Tandon)

Spring 2011

Co-instructed (with Cong Yu) a graduate level seminar on advanced information retrieval. Classes included a combination of lectures and paper discussion. I presented half of the lectures. Class size was 10 students.

Information Retrieval

University of Massachusetts Amherst

Fall 2006

Prepared slides and provided lectures for a graduate-level information retrieval course. Instructor: Professor James Allan.

Applied Information Theory

University of Massachusetts Amherst

Fall 2006

Provided grading support for a graduate-level information theory course. Instructor: Professor Erik Learned-Miller.

Databases

University of Massachusetts Amherst

Spring 2006

Provided teaching assistance for a graduate-level database course. This included grading homeworks, holding office hours, and maintaining a course web page. Instructors: Professors Yanlei Diao and Gerome Miklau.

Tutorials

Fernando Diaz, Andrew Drozdov, To Eun Kim, Alireza Salemi, and Hamed Zamani. Retrieval-Enhanced Machine Learning: Synthesis and Opportunities. In *Proceedings of the 2024 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, 2024

Charles L. A. Clarke, **Fernando Diaz**, and Negar Arabzadeh. Preference-Based Offline Evaluation. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, 2023

- Brian St. Thomas, Praveen Chandar, Christine Hosey, and **Fernando Diaz**. Mixed Method Development of Evaluation Metrics. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021
- Praveen Chandar, **Fernando Diaz**, and Brian St. Thomas. Beyond Accuracy: Grounding Evaluation Metrics for Human-Machine Learning Systems. In *Advances in Neural Information Processing Systems*, 2020
- Michael D Ekstrand, Robin Burke, and **Fernando Diaz**. Fairness and Discrimination in Recommendation and Retrieval. In *Proceedings of the 13th ACM Conference on Recommender Systems*, 2019
- Michael D. Ekstrand, Robin Burke, and **Fernando Diaz**. Fairness and Discrimination in Retrieval and Recommendation. In *Proceedings of the 42Nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2019
- Jean Garcia-Gathright, Christine Hosey, Brian St. Thomas, Ben Carterette, and **Fernando Diaz**. Mixed Methods for Evaluating User Satisfaction. In *Proceedings of the 12th ACM Conference on Recommender Systems*, 2018
- Carlos Castillo, **Fernando Diaz**, and Hemant Purohit. Leveraging Social Media and Web of Data to Assist Crisis Response Coordination. In *SDM '14: Proceedings of the 14th SIAM International Conference on Data Mining*, 2014
- Kira Radinsky, **Fernando Diaz**, Susan Dumais, Milad Shokouhi, Anlei Dong, and Yi Chang. Temporal web dynamics and its application to information retrieval. In *Proceedings of the sixth ACM international conference on Web search and data mining*, 2013
- Fernando Diaz**, Jaime Arguello, and Milad Shokouhi. Integrating and Ranking Aggregated Content on the Web. In *WWW '12: Proceedings of the 21th international conference on World wide web*, 2012
- Fernando Diaz**, Mounia Lalmas, and Milad Shokouhi. From federated to aggregated search. In *SIGIR '10: Proceedings of the 33rd annual international ACM SIGIR conference on Research and development in information retrieval*, 2010

Supervision

Current Students

Alfredo Gomez, with Mona Diab
 Athiya Deviyani
 Jessica Huynh, with Jeffrey Bigham
 Shaily Bhatt
 To Eun Kim

Former Interns

Negar Arabzadeh (2022), PhD Student, University of Waterloo
 Esther Rolf (2021), Postdoctoral Scholar, Harvard University
 Maria Antoniak (2020), Young Investigator, Allen Institute for AI
 Daniel Cohen (2019), Research Scientist, Dataminr
 Ashudeep Singh (2019), Research Scientist, Pinterest
 Jesse Anderton (2017), Research Scientist, Spotify
 Rishabh Mehrotra (2016), Director of Machine Learning, ShareChat
 Cristina Garbacea (2016), PhD Student, University of Michigan
 Chris Kedzie (2015), Senior Researcher, Microsoft Semantic Machines
 David Abel (2015), Research Scientist, DeepMind
 Ioannis Paparrizos (2014), Assistant Professor, Ohio State University
 Matthew Ekstrand-Abeug (2013), Software Engineer, Google

Pavel Metrikov (2013), Data Scientist, Microsoft
Teresa Bracamonte (2013), Developer, Fintual
Qi Guo (2011), Software Engineer, Meta
Jangwon Seo (2010), Staff Software Engineer, Google
Jaime Arguello (2009, 2010), Associate Professor, University of North Carolina Chapel Hill
Ahmed Hassan (2009), Senior Principal Research Manager, Microsoft Research

PhD Examiner

James Route, Carnegie Mellon University, 2025
Jiatong Shi, Carnegie Mellon University, 2025
Till Kletti, Université Grenoble Alpes, 2023
Jesse Anderton, Northeastern University, 2019
Rodrigo Nogueira, New York University, 2019
Matthew Ekstrand-Abeug, Northeastern University, 2017
Maria-Hendrike Peetz, University of Amsterdam, 2015
Jaime Arguello, Carnegie Mellon University, 2011

Patents

Granted

Fernando Diaz, Ryen White, and Qi Guo. Prefetching using dynamic user model to reduce latency. US Patent US12314273B2

Corby Rosset, Bhaskar Mitra, David Hawking, Nick Craswell, **Fernando Diaz**, and Emine Yilmaz. Neural network for search retrieval and ranking, 2023. US Patent US11615149B2

Fernando Diaz, Donald Metzler, and Sihem Amer-Yahia. System for determining and optimizing for relevance in match-making systems, 2016. US Patent US10380158B2

Fernando Diaz. Predicting audience response for scripting, 2015. US Patent US9159031B2

Yi Chang, Zhaohui Zheng, **Fernando Diaz**, and Jing Bai. Cross-market model adaptation with pairwise preference data, 2013. US Patent US8489590B2

Rosie Jones, **Fernando Diaz**, and Ahmed Hassan Awadallah. System and method of geo-based prediction in search result selection, 2013. US Patent US8352466B2

Fernando Diaz. System for integrating content from external corpora into primary search engine results, 2012. US Patent 8,150,874

Rosie Jones and **Fernando Diaz**. System and method for providing temporal search results in response to a search query, 2009. US Patent 7,577,651

Applied

Vanja Josifovski, Evgeniy Gabrilovich, Bo Pang, **Fernando Diaz**, and Jangwon Seo. Method or system for identifying website link suggestions, 2013. US Patent Application US20130173568A1

Elad Yom-Tov and **Fernando Diaz**. Using user's social connection and information in web searching, 2012. US Patent Application US20120295633A1

Recognition

Awards

SIGIR Academy, 2024
CSCW Honorable Mention, 2023
CIKM Best Paper Nomination, 2020
CSCW Honorable Mention, 2020
British Computer Society Karen Spärck Jones Award, 2017
Finalist for Paul E. Green Award, Journal of Marketing Research, 2015
ISCRAM Best Paper Award, 2013
SIGIR Best Paper Nomination, 2011
ECIR Best Student Paper Award, 2011
SIGIR Best Paper Award, 2009
WSDM Best Paper Award, 2009

Fellowships and Chairs

Canada CIFAR AI Chair, 2019
UCLA Institute for Pure and Mathematics Postdoctoral Fellow, 2006
National Science Foundation Summer Institute in Japan, 2001
University of Massachusetts Opportunity Fellowship, 2000

Service

Organizational Leadership

Dagstuhl Seminar: Towards a Multidisciplinary Vision for Culturally Inclusive Generative AI, Co-Organizer, 2025
SIGIR, General Co-Chair, 2021
SIGIR Workshop on Retrieval-Enhanced Machine Learning, Co-Organizer, 2023
TREC Tip of the Tongue Track, Co-Organizer, 2023-2024
TREC Fair Ranking Track, Co-Organizer, 2019-2020
CIFAR Workshop on AI and the Curation of Culture, Co-Organizer, 2019
Strategic Workshop on Information Retrieval in Lorne, Co-Organizer, 2018, 2024
NeurIPS Workshop on Cultures of AI and AI for Culture, Co-Organizer, 2022
CVPR Workshop on Ethical Considerations in Creative applications of Computer Vision, Co-Organizer, 2022
NeurIPS Workshop on Algorithmic Fairness through the Lens of Causality and Interpretability, Co-Organizer, 2020
Workshop on Social Web for Disaster Management, Co-Organizer, 2015-2016
WSDM, General Co-Chair, 2014
TREC Web Track, Co-Organizer, 2013-2014
TREC Temporal Summarization Track, Co-Organizer, 2013-2014
NTCIR Recipe Search Track, Co-Organizer, 2014
SIGIR Workshop on Time-aware Information Access, Co-Organizer, 2012-2014
SIGIR Workshop on Social Web Search and Mining: Analysis of User Generated Content Under Crisis, Co-Organizer, 2011

Organizational Support

ACM FAccT Steering Committee, Member, 2018-present
ACM Ethics & Plagiarism Committee, Member, 2022-present
ACM SIGIR, Awards Chair, 2014-2016

Technical Expertise

FAT*, Program Co-Chair, 2019
SIGIR, Perspectives Track Chair, 2021
SIGIR Forum, Editor, 2012-2014
WWW, Track Chair, 2017
CIKM, Vice Chair, 2009

SIGIR, Senior Program Committee Member, 2011-2020
WSDM, Senior Program Committee Member, 2012-2018, 2023
CIKM, Senior Program Committee Member, 2012
NeurIPS, Ethics Reviewer, 2021-2022
NeurIPS, Datasets and Benchmarks Reviewer, 2022
SIGIR, Perspectives Program Committee Member, 2022
SIGIR, Program Committee Member, 2008-2010
CIKM, Program Committee Member, 2010-2011, 2014
WSDM, Program Committee Member, 2010-2011
WWW, Program Committee Member, 2011-2014

Skills

Extensive programming C/C++ experience on Unix platforms. Moderate experience with audio programming on iOS. Extensive experience with production machine learning environments, including feature generation, model training, and evaluation. Experience with pytorch, GNU Scientific Library, Matlab/Octave, R, Python, Perl, Java, and scientific computing in a clustered environment. English and Spanish fluency. Some French.

References

Available on request.