

# Praveen C. Ravichandran

*curriculum vitae*

1150 N Fayette Street  
Alexandria, VA - 22314

☎ +1 (201) 314 6682

✉ [praveenrchandar@gmail.com](mailto:praveenrchandar@gmail.com)

📄 [pchandar.github.io](https://github.com/pchandar)

8+ years of industry experience in Search, Recommendations, Large-scale Experimentation and Metric Development. Strong emphasis and strength in translating state-of-art research into practice, leading to highly impactful product outcomes. Led projects that spanned across missions and several teams; worked with data scientists, machine learning specialists, researchers, backend, and data engineers. Helped extensively in the hiring process to build research labs and advocated for inclusivity in hiring. Passionate about learning new technologies and efficient at prototyping ambitious ideas that help bootstrap innovative products.

At Spotify, I've led various projects that include development and implementation of processes to improve machine learning experimentation practices across several teams, high-level design of company-level ML infrastructure technologies such as Spotify Kubeflow, and contributions to the company-wide ML Tech Strategy. Previously, I've worked on developing search and conversational AI solutions for meeting scheduling (at X.AI), health care (at Columbia). At IBM, I've focused on the development and implementation of scalable interactive conversational AI agents with a focus on the human-compute-interaction aspects.

I've published research papers on different topics, including search, recommendations, experimentation, chatbots and health care, with 40+ research papers at top conferences, such as WWW, SIGIR, WSDM, KDD, NeurIPS, and served as PC/SPC in several conferences.

---

## Professional Experience

Aug '17 – Present **Staff Research Scientist & Research Lead, Spotify, New York, NY.**

At Spotify, I lead various projects on translating state-of-art research to production. As a tech lead, I work on integrating and expanding Podz (an AI-based Podcast Startup acquired by Spotify) into the company's ecosystem, a core part of the company's podcast strategy. I lead the recommendation efforts, implementing neural recommender models for podcast discovery.

I've also led cross-mission initiatives, working with 8 different ML teams to develop ML best practices within the company. As research lead, I worked with research scientists, data scientists, and engineers to implement causal algorithms to estimate the effect of podcast engagement on financial metrics; outputs from this work were included in Spotify's earnings report in 2020.

As a researcher and engineer, I developed novel off-policy estimation algorithms to improve offline evaluation for recommendation & search systems. I prototyped and tested the latest ML Ops technologies to improve ML Infrastructure at the company. Developed mixed methods for metric development for search and deployed new & effective metric in production.

- Apr '16 – Aug '17 **Research Staff Member**, *IBM Research*, Cambridge, MA.  
At IBM Research AI, my work focused on developing new methods to instantiate interactive conversational AI agents for various business applications. My role involved: (1) building and deploying domain-specific conversational IR systems for HR domains, (2) conducting user studies and controlled experiments to better understand user needs and evaluate conversational systems, and (3) closely working with engineering teams to productionize research prototypes. The work resulted in patents, publications at top conferences, as well as productization of the prototype for instantiating serverless AI agents as a service within IBM Bluemix Cloud.
- Apr '15 – Apr '16 **Data Scientist**, *X.AI Inc.*, New York, NY.  
At X.AI, I led a small team of data scientists and engineers to build NLP components of a dialogue system. Our team developed temporal expression extractors, location and person entity detectors, and email parsers. My role involved writing production code, advising and mentoring junior data scientists, effectively bridging communication gaps between the data science and engineering teams, and helping hire new talent.
- Sept '14 – Apr '15 **Postdoctoral Research Scientist**, *Columbia University*, New York, NY.  
At Columbia, I built dialog systems that enabled easier access of clinical data to researchers. Specifically, I built a framework to analyze informational needs of clinical researchers and developed ML algorithms to enrich medical terminologies.  
Mentor: Dr. Chunhua Weng
- Sept '13 – May '14 **Teaching Assistant**, *University of Delaware*, Newark, DE.
- Feb '09 – May '14 **Research Assistant**, *University of Delaware*, Newark, DE.  
As part of my dissertation, I worked on research problems relating to novelty and diversity in search results. I conducted studies to better understand user needs that lead to the development of a preference framework for novelty evaluation.
- Summer 2013 **Research Intern**, *IBM T.J. Watson Research Center*, Yorktown Heights, NY.  
I worked on methods to evaluate ranking algorithms that retrieve and cluster answer snippets for a given query. I conducted experiments to compare a novel pairwise annotation tool to traditional annotation methods for the BOLT IR project.  
Mentor: Dr. Hema Raghavan
- Summer 2010 **Search Relevance Intern**, *OneRiot Inc (now Walmart Labs)*, Palo Alto, CA.  
As an intern, I worked on building real-time search engines and developed tools to evaluate ranking algorithm that take into account the temporal aspects of relevance.

---

## Education

- Aug 2014 **Ph.D., Computer Science**, *University of Delaware*, Newark, DE.  
*Thesis*: Novelty and Diversity in Search Results  
*Advisor*: Dr. Benjamin A Carterette
- Aug 2010 **M.S., Computer Science**, *University of Delaware*, Newark, DE.
- May 2008 **Bachelor of Engineering in Computer Science**, *Anna University*, Chennai, India.  
*First Class with Distinction*

---

## Recent Invited Talks

- NVIDIA, 2021. A Scalable Experimentation Framework using NVTabular

- RMIT, 2021. Mixed Methods for Metric Development at Spotify
- Peloton, 2021. Metrics & Evaluation for Search and Recommendation

## Professional Activities

### Committees and Reviewing

- *Sr PC Member*: WWW (2020-2021).
- *PC Member*: TWEB, TiOS, AIRS 13, KDIR (16-17), ECIR (14-16), SIGIR (16-21), CIKM (14-20), WSDM (18-21), WWW (15-21), KDD (18-21).

## Research

### Research Interests

Metrics & Experimentation, Information Retrieval, Natural Language Processing, Machine Learning

### Tutorials

- KDD '21 ***Mixed Method Development of Evaluation Metrics*** B St.Thomas, P Chandar, C Hosey, F Diaz (2021)
- NeurIPS '20 ***Beyond Accuracy: Grounding Evaluation Metrics for Human-Machine Learning Systems*** P Chandar, B St.Thomas, F Diaz (2020)

### Peer-Reviewed Publications

- WWW '22 ***Estimating Long-Term Engagement in Online Experiments using Survival Models*** P Chandar, B St. Thomas, L Maystre, V Pappu, R Sanchis Ojeda, T Wu, B Carterette, M Lalmas, T Jebara. (2022)
- WWW '22 ***Choice of Implicit Signal Matters: Accounting for User Aspirations in Podcast Recommendations*** Z Nazari, P Chandar, G Fazelnia, C Edwards, B Carterette, M Lalmas. (2022)
- KDD '21 ***Neural Instant Search for Music and Podcast*** H Hashemi, A Pappu, M Tian, P Chandar, M Lalmas, B Carterette. (2021)
- EC '20 ***The engagement-diversity connection: Evidence from a field experiment on spotify*** D Holtz, B Carterette, P Chandar, Z Nazari, H Cramer, S Aral. (2020)
- WISE '20 ***Methods for individual treatment assignment: An application and comparison for playlist generation*** F Provost, J Anderton, B Carterette, P Chandar. (2020)
- NeurIPS '20 ***Model selection for production system via automated online experiments*** Z Dai, P Chandar, G Fazelnia, B Carterette, M Lalmas. (2020)
- UAI '20 ***Stochastic Variational Inference for Dynamic Correlated Topic Models*** F Tomasi, P Chandar, G Levy-Fix, M Lalmas-Roelleke, Z Dai. (2020)
- KDD '20 ***Counterfactual evaluation of slate recommendations with sequential reward interactions*** P. Chandar, J McInerney, B Brost, R Mehrotra, B Carterette. (2020)
- WWW '20 ***Do podcasts and music compete with one another? Understanding users' audio streaming habits*** A Li, A Wang, Z Nazari, P Chandar, B Carterette. (2020)
- SIGIR '19 ***Developing evaluation metrics for instant search using mixed methods*** P Chandar, J Garcia-Gathright, C Hosey, B St. Thomas, J Thom. (2019)

- WWW '19 ***Search mindsets: Understanding focused and non-focused information seeking in music search*** A Li, J Thom, P Chandar, C Hosey, BS Thomas, J Garcia-Gathright. (2019)
- WSDM '19 ***Offline evaluation to make decisions about playlist recommendation algorithms*** A Gruson, P Chandar, C Charbuillet, J McInerney, S Hansen, D Tardieu, B. Carterette. (2019)
- CIKM '18 ***Estimating clickthrough bias in the cascade model*** P. Chandar, B. Carterette. (2018)
- SIGIR '18 ***An information retrieval framework for contextual suggestion based on heterogeneous information network embeddings*** D Seyler, P Chandar, M Davis. (2018)
- SIGIR '18 ***Offline comparative evaluation with incremental, minimally-invasive online feedback*** B. Carterette, P. Chandar (2018)
- CHI '18 ***All Work and no Play? Conversations with a Question-and-Answer Chatbot in the Wild*** Q.V. Liao, M. Mas-ud Hussain, P. Chandar, M. Davis, Y. Khazaeni, M. Crasso, D. Wang, M. Muller, N. S. Shami, W. Geyer. (2018)
- INTERACT '17 ***Leveraging Conversational Agent to assists New Hires during Onboarding*** P. Chandar, Y. Khazaeni, M. Davis, M. Muller, M. Crasso, Q.V. Liao, N. S. Shami, W. Geyer. (2017)
- workshop @ IJCAI '16 ***Using an AI Agent and Coordinated Expert Sourcing to Construct Content for a Dialog System*** M. Davis, W. Geyer, Y. Khazaeni, M. Crasso, P. Chandar and D. Levitin. (2016)
- JMI '16 ***A Data-driven Concept Schema for Defining Clinical Research Data Needs*** Hruby G.W., Hoxha J., Chandar P., Mendonca E., Hanauer D., Weng C. (2016) Vol-91 Issue-1 pp. 1–9 )
- JBIR '16 ***DREAM: Classification Scheme for Dialog Acts in Clinical Research Query Mediation*** Hoxha J, Chandar P, He Z, J. Cimino, D. Hanauer, Weng C. (2016) Vol-59 Issue-1 pp. 89–101)
- AMIA '15 ***Simulation-based Evaluation of the Generalizability Index for Study Traits*** Z. He, P. Chandar, P. Ryan, C. Weng. *Proceedings of AMIA 2015 Annual Symposium (2015)*, pp.593–602
- AMIA '15 ***Similarity-based Recommendation of New Concepts to a Terminology*** P. Chandar, A. Yaman, J. Hoxha, Z. He, C. Weng. *Proceedings of AMIA 2015 Annual Symposium (2015)*, pp.386–395
- AMIA '15 ***What Are Frequent Data Requests from Researchers? A Conceptual Model of Researchers' EHR Data Needs for Comparative Effectiveness Research Podium abstract*** G. Hruby, P. Chandar, J. Hoxha, E. Mandonca, D. Hanauer, C. Weng. *Proceedings of AMIA 2015 Annual Symposium (2015)*, Podium Abstract
- SIGIR '15 ***Document Comprehensiveness and User Preferences in Novelty Search Tasks*** A. Bah, P. Chandar, B. Carterette. *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval (2015)*, pp. 735–738
- SIGIR '13 ***Preference Based Evaluation Measures for Novelty and Diversity*** P. Chandar, B. Carterette. In *Proceedings of the 36th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. (2013)*, pp. 413–422.
- SIGIR '13 ***Document Features Predicting Assessor Disagreement*** P. Chandar, W. Webber, B. Carterette. In *Proceedings of the 36th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (2013)*, pp. 745–748.

- CIKM '12 **Alternative Assessor Disagreement and Retrieval Depth** W. Webber, P. Chandar, B. Carterette. In *Proceedings of the 21st ACM International Conference on Information and Knowledge Management (2012)*, pp. 125–134.
- SIGIR '12 **Using Preference Judgments for Novel Document Retrieval** P. Chandar, B. Carterette. In *Proceedings of the 35th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (2012)*, pp. 861–870.
- SIGIR '12 **Using PageRank to Infer User Preferences** P. Chandar, B. Carterette. In *Proceedings of the 35th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (2012)*, pp. 1167–1168.
- workshop @ WSDM '12 **What Qualities Do Users Prefer in Diversity Rankings?** P. Chandar, B. Carterette. *Workshop on Diversity in Document Retrieval at the fifth ACM WSDM Conference (2012)*.
- workshop @ ECIR '11 **Analysis of Various Evaluation Measures for Diversity** P. Chandar, B. Carterette. *Workshop on Diversity in Document Retrieval at 33rd European Conference on Information Retrieval (ECIR) (2011)*.
- SIGIR '10 **Diversification of Search Results using Webgraphs** P. Chandar, B. Carterette. In *Proceedings of the 33rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (2010)*, pp. 869–870.
- CIKM '09 **Probabilistic Models for Facet Retrieval** B. Carterette, P. Chandar. In *Proceedings of the 18th ACM International Conference on Information and Knowledge Management 2009*, pp. 1287–1296.

### Technical Reports and Other Publications

- NTCIR '14 **Udel @ NTCIR-11 IMine Track** A. Bah, P. Chandar, B. Carterette. *Proceedings of the 11th NTCIR Conference (2014)*.
- TREC '13 **University of Delaware at TREC 2013** A. Bah, P. Chandar, K. Sabhnani, M. Zengin, B. Carterette. *Proceedings of the 22nd Text Retrieval Conference (2013)*.
- TREC '12 **University of Delaware at TREC 2012** A. Bah, P. Chandar, N. Kumar, A. Rao, D. Zhu, B. Carterette. *Proceedings of the 21st Text Retrieval Conference (2012)*.
- TREC '11 **Implicit Feedback and Document Filtering for IR Over Query Sessions** B. Carterette, P. Chandar. *Proceedings of the 20th Text Retrieval Conference (2011)*.
- TREC '10 **Sessions, Diversity, and Ad Hoc Retrieval** B. Carterette, P. Chandar. *Proceedings of the 19th Text Retrieval Conference (2010)*.
- TREC '09 **Ad Hoc and Diversity Retrieval at the University of Delaware** P. Chandar, A. Kailasam, D. Muppaneni, L. Thota, B. Carterette. *Proceedings of the 18th Text Retrieval Conference (2009)*.
- TREC '09 **Minimal Test Collections for Relevance Feedback** B. Carterette, P. Chandar, A. Kailasam, D. Muppaneni, L. Thota. *Proceedings of the 18th Text Retrieval Conference (2009)*.

## Teaching

### Teaching Assistant, University of Delaware

#### CISC 106 **Introductory Computer Science for Engineers**

- Instructor: Debra Yarrington
- Assisted in creating lab and homework assignments.

- Conducted lab sessions and graded assignments.
- CISC 672/471 ***Compiler Design***
- Instructor: Prof. Lori Pollock
  - Lectured on the following topics: Bottom-up parsing, and Abstract syntax trees
  - Assisted in construction of lab projects and homework assignments.
  - Held office hours, responding to student concerns on course material

---

## Proficiencies

Programming	Python, Java, Scala, C++, R
ML Frameworks	PyTorch, Tensorflow/Keras, NumPyro, EconML
MLOps	TFX, Kubeflow Pipelines, NVTabular

---

## References

- Available upon request