



## Paul N. Bennett

Senior Researcher, CLUES  
Microsoft Research  
One Microsoft Way  
Redmond, WA 98052

Cell: +1 (425) 301-1231

<http://research.microsoft.com/~paubenpbennett+cv@cs.cmu.edu>

---

## Research Interests

*Machine learning, information retrieval, data mining, text analysis*

I am a data mining and machine learning researcher situated as a core member of the information retrieval community working on web scale challenges and intelligent virtual assistants. I advance fundamental techniques of machine learning and information retrieval: in particular ensemble methods and the combination of information sources, calibration, consensus methods for noisy supervision labels, active learning and evaluation, and supervised classification (with an emphasis on hierarchical classification) and ranking. I then apply these methods and other machine learning methods to mine probabilistic representations of contextual information signals from text, interaction streams, and other information sources (e.g. demographic, location, social) for interpretive purposes to aid data understanding and user behavior and improve personalized user models. Using these insights and mined representations, I improve user-facing intelligent applications via such techniques as the incorporation of contextual ranking and recommendation models which can operate at scale while incorporating the uncertain information in the underlying contextual signals. More generally, I am interested in the development, improvement, and analysis of machine learning and data mining methods incorporated as components of real-world intelligent personal systems, large-scale adaptive systems, or information retrieval systems.

## Education

- **Carnegie Mellon University**: August 1998--May 2006
  - **Ph.D. Computer Science**, Computer Science Department, May 2006  
Advisor: [Jaime Carbonell](#), Co-advisor: [John Lafferty](#)
  - **M.S. Computer Science**, Computer Science Department, August 2000
- **University of Texas at Austin**: August 1993--May 1998
  - **B.S. in Computer Sciences** with Special Honors in Computer Sciences 5/98; cum laude
  - **B.A. in Philosophy** 12/97; **Secondary Major certification in Plan II Honors**; cum laude
  - GPA information available upon request.
- Pre-university information available upon request.

## Publications

### Journal Papers

- **[Understanding Intrinsic Diversity in Web Search: Improving Whole-Session Relevance](#)**  
Karthik Raman, Paul N. Bennett, Kevyn Collins-Thompson  
In *ACM Transactions on Information Systems*. vol. 32 (3), 20:1-20:25. 2014.
- **[Learning to Rank using an Ensemble of Lambda-Gradient Models](#)**  
Christopher J.C. Burges, Krysta M. Svore, Paul N. Bennett, Andrzej Pastusiak, and Qiang Wu  
In *Journal of Machine Learning Research: Workshop and Conference Proceedings, Yahoo! Learning to Rank Challenge*. vol. 14, pp. 25-35, Journal of Machine Learning Research. 2011.
- **[The Combination of Text Classifiers Using Reliability Indicators](#)**  
Paul N. Bennett, Susan T. Dumais, Eric Horvitz  
*Information Retrieval*. Kluwer Academic Publishers.  
January, 2005, 8(1):67-100.  
(This paper revises and extends work originally presented in [Bennett, Dumais, & Horvitz; SIGIR 2002].)

### Conference Papers

- Rigorously Reviewed Conferences (Acceptance Rate  $\leq$  25%)
  - **[Characterizing and Predicting Enterprise Email Reply Behavior](#)**  
Liu Yang, Susan, T. Dumais, Paul N. Bennett, and Ahmed Hassan Awadallah  
To appear in *Proceedings of the 40th Annual ACM SIGIR Conference (SIGIR 2017)*. Tokyo, Japan.  
August 2017.

- [\*\*Managing Uncertainty in Time Expressions for Virtual Assistants\*\*](#)  
Xin Rong, Adam Fourney, Robin N. Brewer, Meredith Ringel Morris, Paul N. Bennett  
In *Proceedings of the 35th Annual CHI Conference on Human Factors in Computing Systems* (CHI '17). Denver, Colorado. May, 2017. (*Best Paper Honorable Mention*)
- [\*\*Algorithms for Active Classifier Selection: Maximizing Recall with Precision Constraints\*\*](#)  
Paul N. Bennett, David M. Chickering, Christopher Meek, Xiaojin Zhu  
In *Proceedings of the 10th ACM International Conference on Web Search and Data Mining* (WSDM '17). Cambridge, UK. February, 2017.
- [\*\*Analyzing and Predicting Task Reminders\*\*](#)  
David Graus, Paul N. Bennett, Ryen White, Eric Horvitz  
In *Proceedings of the 24th ACM Conference on User Modeling, Adaptation, and Personalization* (UMAP '16). Halifax, Canada. July 2016. (*Best Student Paper*)
- [\*\*Using Shortlists to Support Decision Making and Improve Recommender System Performance\*\*](#)  
Tobias Schnabel, Paul N. Bennett, Susan T. Dumais, Thorsten Joachims  
In *Proceedings of the 25th International World Wide Web Conference* (WWW '16). Montreal, Canada. April 2016.
- [\*\*Events and Controversies: Influences of a Shocking News Event on Information Seeking\*\*](#)  
Danai Koutra, Paul N. Bennett, Eric Horvitz  
In *Proceedings of the 24th Annual International World Wide Web Conference* (WWW '15). Florence, Italy. May 2015.  
(This paper revises and extends work in TAIA Workshop at SIGIR 2012 and arXiv:1405.1486.)
- [\*\*Query Suggestion and Data Fusion in Contextual Disambiguation\*\*](#)  
Milad Shokouhi, Marc Sloan, Paul N. Bennett, Kevyn Collins-Thompson, Siranush Sarkizova  
In *Proceedings of the 24th Annual International World Wide Web Conference* (WWW '15). Florence, Italy. May 2015.
- [\*\*Overcoming Relational Learning Biases to Accurately Predict Preferences in Large Scale Networks\*\*](#)  
Joseph J. Pfeiffer III, Jennifer Neville, Paul N. Bennett  
In *Proceedings of the 24th Annual International World Wide Web Conference* (WWW '15). Florence, Italy. May 2015.
- [\*\*Composite Likelihood Data Augmentation for Within-Network Statistical Relational Learning\*\*](#)  
Joseph J. Pfeiffer III, Jennifer Neville, and Paul N. Bennett  
In *Proceedings of the 14th IEEE International Conference on Data Mining* (ICDM 2014). Shenzhen, China. December, 2014.
- [\*\*Active Exploration in Networks: Using Probabilistic Relationships for Learning and Inference\*\*](#)  
Joseph J. Pfeiffer III, Jennifer Neville, and Paul N. Bennett  
In *Proceedings of the 23rd ACM International Conference on Information and Knowledge Management* (CIKM 2014). Shanghai, China. October, 2014.
- [\*\*Toward Whole-Session Relevance: Exploring Intrinsic Diversity in Web Search\*\*](#)  
Karthik Raman, Paul N. Bennett, and Kevyn Collins-Thompson  
In *Proceedings of the 36th Annual ACM SIGIR Conference* (SIGIR 2013). Dublin, Ireland. July-August 2013. (*Best Student Paper Award*)
- [\*\*Fighting Search Engine Amnesia: Reranking Repeated Results\*\*](#)  
Milad Shokouhi, Ryen White, Paul N. Bennett, Filip Radlinski  
In *Proceedings of the 36th Annual ACM SIGIR Conference* (SIGIR 2013). Dublin, Ireland. July-August 2013.
- [\*\*Pairwise Ranking Aggregation in a Crowdsourced Setting\*\*](#)  
Xi Chen, Paul N. Bennett, Kevyn Collins-Thompson, and Eric Horvitz  
In *Proceedings of the 6th ACM International Conference on Web Search and Data Mining* (WSDM '13). Rome, Italy. February 2013.
- [\*\*Predicting Content Change on the Web\*\*](#)  
Kira Radinsky and Paul N. Bennett  
In *Proceedings of the 6th ACM International Conference on Web Search and Data Mining* (WSDM '13). Rome, Italy. February 2013.
- [\*\*Personalizing Atypical Web Search Sessions\*\*](#)  
Carsten Eickhoff, Kevyn Collins-Thompson, Paul N. Bennett, and Susan Dumais  
In *Proceedings of the 6th ACM International Conference on Web Search and Data Mining* (WSDM '13). Rome, Italy. February 2013.
- [\*\*Modeling the Impact of Short- and Long-Term Behavior on Search Personalization\*\*](#)  
P.N. Bennett, R.W. White, W. Chu, S.T. Dumais, P. Bailey, F. Borisyuk, and X. Cui  
In *Proceedings of the 35th Annual ACM SIGIR Conference* (SIGIR 2012). Portland, WA. August 2012.

- [\*\*Robust Ranking Models via Risk-Sensitive Optimization\*\*](#)  
L. Wang, P.N. Bennett, K. Collins-Thompson  
In *Proceedings of the 35th Annual ACM SIGIR Conference (SIGIR 2012)*. Portland, WA. August 2012.  
(Honorable Mention Best Paper)
- [\*\*Search, Interrupted: Understanding and Predicting Search Task Continuation\*\*](#)  
E. Agichtein, R.W. White, S.T. Dumais, and P.N. Bennett  
In *Proceedings of the 35th Annual ACM SIGIR Conference (SIGIR 2012)*. Portland, WA. August 2012.
- [\*\*Probabilistic Models for Personalizing Web Search\*\*](#)  
David Sontag, Kevyn Collins-Thompson, Paul N. Bennett, Ryen White, Susan Dumais, and Bodo von Billerbeck  
In *Proceedings of the 5th ACM International Conference on Web Search and Data Mining (WSDM '12)*. Seattle, Washington. February 2012.
- [\*\*Characterizing Web Content, User Interests, and Search Behavior by Reading Level and Topic\*\*](#)  
Jinyoung Kim, Kevyn Collins-Thompson, Paul N. Bennett, and Susan Dumais  
In *Proceedings of the 5th ACM International Conference on Web Search and Data Mining (WSDM '12)*. Seattle, Washington. February 2012.
- [\*\*Personalizing Web Search Results by Reading Level\*\*](#)  
Kevyn Collins-Thompson, Paul N. Bennett, Ryen White, Sebastian de la Chica, and David Sontag  
In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management (CIKM '11)*. Glasgow, Scotland. October 2011.
- [\*\*Inferring and Using Location Metadata to Personalize Web Search\*\*](#)  
Paul N. Bennett, Filip Radlinski, Ryen W. White, and Emine Yilmaz  
In *Proceedings of the 34th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2011)*. Beijing, China. July 2011.
- [\*\*Modeling and Analysis of Cross-Session Search Tasks\*\*](#)  
Alex Kotov, Paul N. Bennett, Ryen W. White, Susan Dumais, and Jaime Teevan  
In *Proceedings of the 34th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2011)*. Beijing, China. July 2011.
- [\*\*Fractional Similarity: Cross-lingual Feature Selection for Search\*\*](#)  
Jagadeesh Jagarlamudi and Paul N. Bennett  
In *Proceedings of the 33rd Annual European Conference on Information Retrieval (ECIR 2011)*. Dublin, Ireland. April 2011.
- [\*\*Detecting Duplicate Web Documents using Clickthrough Data\*\*](#)  
F. Radlinski, P.N. Bennett, E. Yilmaz  
In *Proceedings of the 4th ACM International Conference on Web Search and Data Mining (WSDM '11)*. Hong Kong. Feb 2011.
- [\*\*Predicting Short-term Interests Using Activity-Based Search Contexts\*\*](#)  
R. White, P.N. Bennett, S. Dumais  
In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management (CIKM '10)*. Toronto, Canada. Oct 2010.
- [\*\*Classification-Enhanced Ranking\*\*](#)  
P.N. Bennett, K. Svore, S. Dumais  
In *Proceedings of the 19th Annual International World Wide Web Conference (WWW '10)*. Raleigh, NC. April 2010.
- [\*\*Predicting Query Performance via Classification\*\*](#)  
K. Collins-Thompson, P.N. Bennett  
In *Proceedings of the 32nd Annual European Conference on Information Retrieval (ECIR 2010)*. Milton Keynes, UK. March 2010.
- [\*\*Refined Experts: Improving Classification in Large Taxonomies\*\*](#)  
P.N. Bennett, N. Nguyen  
In *Proceedings of the 32nd Annual ACM SIGIR Conference (SIGIR 2009)*. Boston, MA. July 2009.
- [\*\*Learning Consensus Opinion: Mining Data from a Labeling Game\*\*](#)  
Paul N. Bennett, David Maxwell Chickering, Anton Mityagin  
In *Proceedings of the 18th Annual International World Wide Web (WWW) Conference 2009*, Madrid, Spain, April, 2009.
- [\*\*Here Or There: Preference Judgments for Relevance\*\*](#)  
Ben Carterette, Paul N. Bennett, David Maxwell Chickering, Susan T. Dumais  
In *Proceedings of the European Conference on Information Retrieval (ECIR) 2008*, Glasgow, Scotland, March-April, 2008.

- [\*\*Neighborhood-Based Local Sensitivity\*\*](#)  
Paul N. Bennett  
In *Proceedings of the European Conference on Machine Learning (ECML) 2007*, Warsaw, Poland, September, 2007.
- [\*\*Dual Strategy Active Learning\*\*](#)  
Pinar Donmez, Jaime G. Carbonell, Paul N. Bennett  
In *Proceedings of the European Conference on Machine Learning (ECML) 2007*, Warsaw, Poland, September, 2007.
- [\*\*Combining Probability-Based Rankers for Action-Item Detection\*\*](#)  
Paul N. Bennett and Jaime G. Carbonell  
In *Proceedings of the Human Language Technologies — North American ACL 2007 Conference*, Rochester, New York, April, 2007.
- [\*\*Using Asymmetric Distributions to Improve Text Classifier Probability Estimates\*\*](#)  
Paul N. Bennett  
*Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Toronto, Canada, July-August 2003. ACM Press.  
(This paper revises work that originally appeared as CMU-CS-02-126.)
- [\*\*Probabilistic Combination of Text Classifiers Using Reliability Indicators: Models and Results\*\*](#)  
Paul N. Bennett, Susan T. Dumais, Eric Horvitz  
*Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Tampere, Finland, August 2002. ACM Press.  
(An extended and revised version of this work appears in *Information Retrieval*, 2005.)
- Other Conferences (Acceptance Rate > 25%)
  - [\*\*Self-Es: The Role of Emails-to-Self in Personal Information Management\*\*](#)  
Horatiu Bota, Paul N. Bennett, Ahmed Hassan Awadallah, Susan T. Dumais  
In *Proceedings of the 2<sup>nd</sup> ACM Conference on Human Information Interaction and Retrieval (CHIIR '17)*. Oslo, Norway. March, 2017.
  - [\*\*Mining Tasks from the Web Anchor Text Graph\*\*](#)  
Paul N. Bennett, Ryen White  
In *Proceedings of the 24th Text REtrieval Conference (TREC 2015)*. Gaithersburg, Maryland. January, 2016. Related talk slides: [\*\*Mining Tasks from the Web Anchor Text Graph\*\*](#).
  - [\*\*Implicit Preference Labels for Learning Highly Selective Personalized Rankers\*\*](#)  
Paul N. Bennett, Milad Shokouhi, Rich Caruana  
In *Proceedings of the ACM SIGIR International Conference on the Theory of Information Retrieval (ICTIR '15)*. Northampton, MA, USA. September 2015.
  - [\*\*MSRC at TREC 2011 Crowdsourcing Track\*\*](#)  
Paul N. Bennett, Ece Kamar, Gabriella Kazai  
In *Proceedings of the Twentieth Text REtrieval Conference (TREC 2011)*. Gaithersburg, Maryland. November 2011. Related talk slides: [\*\*Consensus Relevance with Topic and Worker Conditional Models\*\*](#).
  - [\*\*Reducing Boundary Friction Using Translation-Fragment Overlap\*\*](#)  
Ralf Brown, Rebecca Hutchinson, Paul N. Bennett, Jaime Carbonell, Peter Jansen  
*Proceedings of the Machine Translation Summit IX*, New Orleans, U.S.A., September 2003.  
(Shares significant overlap in content with publication listing for CMU-CS-03-138, 2003.)

## Other Peer Reviewed Publications

### Poster and Short Papers

- [\*\*Building Effective Query Classifiers: A Case Study in Self-Harm Intent Detection\*\*](#)  
Ashiqur KhudaBukhsh, Paul N. Bennett, Ryen W. White  
Poster-Paper in *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM 2015)*. Melbourne, Australia. October, 2015.
- [\*\*Modeling Website Topic Cohesion at Scale to Improve Webpage Classification\*\*](#)  
Dhivya Eswaran, Paul N. Bennett, Joseph J. Pfeiffer III  
Poster-Paper in *Proceedings of the 38th Annual ACM SIGIR Conference (SIGIR 2015)*. Santiago, Chile. August 2015.
- [\*\*Designing Human-Readable User Profiles for Search Evaluation\*\*](#)  
Carsten Eickhoff, Kevyn Collins-Thompson, Paul N. Bennett, and Susan Dumais  
Short-paper in *Proceedings of the 35th ACM European Conference on Information Retrieval (ECIR 2013)*. Moscow, Russia. March 2013.

- [Active Query Selection for Learning Rankers](#)  
M. Bilgic and P.N. Bennett  
Poster-Paper in *Proceedings of the 35th Annual ACM SIGIR Conference (SIGIR 2012)*. Portland, WA. August 2012.
- [Anticipatory Search: Using Context to Initiate Search](#)  
D.J. Liebling, P.N. Bennett, and R.W. White  
Poster-Paper in *Proceedings of the 35th Annual ACM SIGIR Conference (SIGIR 2012)*. Portland, WA. August 2012.
- [Leveraging Interlingual Classification to Improve Web Search](#)  
Jagadeesh Jagarlamudi, Paul N. Bennett, and Krysta Svore  
In Poster-Track of the *Proceedings of the 21st Annual International World Wide Web Conference (WWW '12)*. Lyon, France. April 2012.
- [The Effects of Choice in Routing Relevance Judgments](#)  
Edith Law, Paul N. Bennett, and Eric Horvitz  
Appears as a Poster-Paper in *Proceedings of the 34th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2011)*. Beijing, China. July 2011.
- [Online Stratified Sampling: Evaluating Classifiers at Web-Scale](#)  
P.N. Bennett, V. Carvalho  
Short-paper in *Proceedings of the 19th ACM International Conference on Information and Knowledge Management (CIKM '10)*. Toronto, Canada. Oct 2010.
- [Estimating Query Performance using Class Predictions.](#)  
K. Collins-Thompson, P.N. Bennett  
Appears as a Poster-Paper in *Proceedings of the 32nd Annual ACM SIGIR Conference (SIGIR 2009)*. Boston, MA. July 2009.
- [Evaluation Measures for Preference Judgments](#)  
Ben Carterette, Paul N. Bennett  
Poster-paper in *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Singapore, July, 2008.
- [Detecting Action-Items in E-mail](#)  
Paul N. Bennett and Jaime Carbonell  
Appears as a Poster-Paper in *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Salvador, Brazil, August 2005.

## Workshops

- [Characterizing Cross-Domain Search Behavior](#)  
Rishabh, Mehrotra, Paul N. Bennett, Susan Dumais, Jaime Teevan.  
In *Proceedings of the Second Workshop on Heterogeneous Information Access at SIGIR 2016 (HIA'16)*. Pisa, Italy. July, 2016.
- [Persona-ization: Searching on Behalf of Others](#)  
Paul N. Bennett, Emre Kiciman.  
In *SIGIR 2015 International Workshop on Social Personalisation & Search (SPS'15)*, August 2015.
- [Influences of a Shocking News Event on Web Browsing](#)  
Danai Koutra, Paul N. Bennett, Eric Horvitz.  
In *SIGIR 2014 Workshop on Temporal, Social and Spatially-aware Information Access (TAIA'14)*, July 2014.  
(This paper was revised and extended in WWW 2015 paper of similar name.)
- [Personalization via Probabilistic Adaptation](#)  
D. Sontag, K. Collins-Thompson, P.N. Bennett, R.W. White, S.T. Dumais, B. Billerbeck  
In *Proceedings of the NIPS 2013 Workshop on Personalization*. Reno, USA.  
(An extended abstract of work that first appeared in WSDM 2012).
- [Combining Active Sampling with Parameter Estimation and Prediction in Single Networks](#)  
J.J. Pfeiffer, J. Neville, P.N. Bennett  
In *Proceedings of the ICML 2013 Workshop on Structured Learning: Inferring Graphs from Structured and Unstructured Inputs (SLG 2013)*. Atlanta, USA. June 2013.
- [Active Sampling of Networks](#)  
J.J. Pfeiffer, J. Neville, P.N. Bennett  
In *Proceedings of the ICML 2012 Workshop on Mining and Learning with Graphs*. Edinburgh, Scotland. July 2012.
- [Picture This: Preferences for Image Search.](#)  
P.N. Bennett, D.M. Chickering, A. Mityagin  
Demo-paper in the *Human Computation 2009 Workshop at KDD '09* (see WWW'09 paper for primary version). Paris, France. June 2009.
- [A Test Collection of Preference Judgments](#)  
Ben Carterette, Paul N. Bennett, Olivier Chapelle  
In *Proceedings of the SIGIR 2008 Beyond Binary Relevance: Preferences, Diversity, and Set-Level Judgments Workshop*, Singapore, July, 2008.

- [\*Feature Representation for Effective Action-Item Detection\*](#)  
Paul N. Bennett and Jaime Carbonell  
*Beyond Bag-of-Words Workshop at the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 63-70, Salvador, Brazil, August 15 - August 19, 2005. ACM Press.  
(This extends work presented as a SIGIR 2005 poster.)
- [\*Inductive Transfer for Text Classification using Generalized Reliability Indicators\*](#)  
Paul N. Bennett, Susan T. Dumais, Eric Horvitz  
*Proceedings of the ICML-2003 Workshop on The Continuum from Labeled to Unlabeled Data*, Washington DC, U.S.A., August 2003.
- [\*Book Recommending Using Text Categorization with Extracted Information\*](#)  
Raymond J. Mooney, Paul N. Bennett and Lorie Roy  
*In the AAAI-98/ICML-98 Workshop on Learning for Text Categorization and the AAAI-98 Workshop on Recommender Systems*, Madison, WI, July 1998.

## Technical Reports and Other Unreviewed Publications

- [\*Overview of the TREC 2014 Web Track\*](#)  
Kevyn Collins-Thompson, Craig Macdonald, Paul Bennett, Fernando Diaz, Ellen M. Voorhees.  
In *Proceedings of the 23rd Text REtrieval Conference*. Feb 2015.
- [\*Events and Controversies: Influences of a Shocking News Event on Information Seeking\*](#)  
Danai Koutra, Paul N. Bennett, Eric Horvitz.  
arXiv:1405.1486, May 2014.  
(This paper was revised and extended in WWW 2015 paper of similar name.)
- [\*Overview of the Special Issue on Contextual Search and Recommendation\*](#)  
Paul N. Bennett, Kevyn Collins-Thompson, Diane Kelly, Ryan W. White, Yi Zhang.  
In *ACM Transactions on Information Systems*. vol. 33 (1). 2015.
- [\*The Emergence of a Peer Relationship in E-mail: A Longitudinal Study of a Student-Advisor Relationship\*](#)  
Paul N. Bennett  
In *Proceedings of JGC60: A Celebration of the Life and Work of Jaime G. Carbonell*. Pittsburgh, PA. April 2014.
- [\*Overview of the TREC 2013 Web Track\*](#)  
Kevyn Collins-Thompson, Paul N. Bennett, Fernando Diaz, Charlie Clarke, Ellen Voorhees  
In *Proceedings of the 22nd Text REtrieval Conference*. Feb 2014.
- [\*Enriching Information Retrieval \(Summary of the SIGIR 2011 Workshop\)\*](#)  
Paul N. Bennett, Khalid El-Arini, Thorsten Joachims, Krysta M. Svore  
In *SIGIR Forum*. pp. 60-65. Dec 2011.
- [\*Redundancy, Diversity, and Interdependent Document Relevance\*](#), a summary of the SIGIR 2009 workshop.  
F. Radlinski, P.N. Bennett, B. Carterette, T. Joachims  
In *SIGIR Forum*. pp. 46-52. Dec 2009.
- [\*A report on the Human Computation Workshop \(HComp 2009\)\*](#)  
P.G. Ipeirotis, R. Chandrasekar, P.N. Bennett, E. Law, M. Chickering, A. Mityagin, F. Provost, L. von Ahn  
In *KDD Explorations*. Vol 11:2 pp. 80-83. Dec 2009.
- [\*Beyond Binary Relevance: Preferences, Diversity, and Set-Level Judgments\*](#), a summary of the SIGIR 2008 workshop.  
P.N. Bennett, B. Carterette, O. Chapelle, T. Joachims  
In *SIGIR Forum*. pp 53-58. Dec 2008.
- [\*Combining Personalized Agents to Improve Content-Based Recommendations\*](#)  
J.M. Adams, P.N. Bennett, and A. Tomasic  
CMU-LTI-07-015, Language Technologies Institute, School of Computer Science, Carnegie Mellon University, December 2007.
- [\*Building Reliable Metaclassifiers for Text Learning\*](#)  
Paul N. Bennett  
CMU-CS-06-121, Ph.D. Thesis, Computer Science Department, School of Computer Science, Carnegie Mellon University, May 2006.
- [\*Maximal Lattice Overlap in Example-Based Machine Translation\*](#)  
Rebecca Hutchinson, Paul N. Bennett, Jaime Carbonell, Peter Jansen, Ralf Brown  
CMU-CS-03-138, Computer Science Department, School of Computer Science, Carnegie Mellon University, June 2003.  
(Also listed as CMU-LTI-03-174. Shares significant overlap in content with publication listing for MT Summit IX, 2003.)

- [\*Using Asymmetric Distributions to Improve Classifier Probabilities: A Comparison of New and Standard Parametric Methods\*](#)  
Paul N. Bennett  
CMU-CS-02-126, Computer Science Department, School of Computer Science, Carnegie Mellon University, April 2002.  
(A revised version of this work appears in SIGIR 2003.)
- [\*Assessing the Calibration of Naive Bayes' Posterior Estimates\*](#)  
Paul N. Bennett  
CMU-CS-00-155, Computer Science Department, School of Computer Science, Carnegie Mellon University, September 2000.
- [\*Text Categorization Through Probabilistic Learning: Applications to Recommender Systems\*](#)  
Paul N. Bennett  
Undergraduate Honor Thesis, Department of Computer Sciences, University of Texas at Austin, May 1998. Also appears as AI TR 98-270.

## Books

- [\*The Guide to Living in Pittsburgh, 9th Edition\*](#)  
Paul N. Bennett and Angela Demke Brown, editors.  
(2000), 299 pages, Pittsburgh, PA: Dec/5, Inc.  
[More information](http://www.cs.cmu.edu/~pghguide) is available for the Guide at <http://www.cs.cmu.edu/~pghguide> .

## Editor

- *Special Issue of ACM TOIS on Contextual Search and Recommendation*  
Paul N. Bennett, Kevyn Collins-Thompson, Diane Kelly, Ryen W. White, Yi Zhang (eds.).  
In *ACM Transactions on Information Systems*. vol. 33 (1). 2015.

## Other Presentations & Panels

- *Events and Controversies: Influences of a Shocking News Event on Information Seeking*. University of Texas at Austin School of Information Colloquium, November 24, 2015.
- *Search from Personal to Social Context: Progress and Challenges*. Keynote talk at the *International Workshop on Social Personalisation & Search* at the 38<sup>th</sup> Annual ACM SIGIR Conference (SIGIR 2015). Santiago, Chile August, 13, 2015.
- *Events and Controversies: Influences of a Shocking News Event on Information Seeking*. Invited talk at Cornell University, Computer Science Department, June 8, 2015.
- *Toward Whole-Session Relevance: Exploring Intrinsic Diversity in Web Search*. Invited talk at Oregon State School of Electrical Engineering and Computer Science Colloquium Series. May 19, 2014.
- *Extreme Classification in Large Taxonomies*. Invited talk at NIPS 2013 Workshop on *Xtreme Classification*.
- CIKM 2013 Panelist on *Channeling the Deluge: Research Challenges for Big Data and Information Systems*.
- *Mining, Understanding, and Modeling Context in Web Search*. Invited talk at Syracuse University School of Information, Brown Bag Seminar, September 2013.
- *Mining and Learning from Context in Web Search*. Invited talk at Cornell University, Artificial Intelligence Seminar, September 2013.
- *Toward Whole-Session Relevance: Exploring Intrinsic Diversity in Web Search*. Invited Talk at University of Illinois at Urbana-Champaign, Yahoo!-DAIS seminar. September 2013.
- Invited Panelist at SIGIR 2013 Workshop on *Benchmarking Adaptive Retrieval and Recommender Systems*. August, 2013
- *Mining and Learning from Context in Web Search*. Invited talk at Carnegie Mellon University School of Computer Science; University of Michigan School of Information; NYU Stern School of Business; University of Texas McCombs School of Business. Feb-April 2013.
- *Mining and Learning from Context in Search*. Keynote talk at the *Context-awareness in Retrieval and Recommendation Workshop* at the 6<sup>th</sup> ACM International Conference on Web Search and Data Mining (WSDM 2013). Feb 2013.
- *Proprietary Data in Research: Public resources and questions of reproducibility*. Invited talk at the *Web Search Click Data Workshop* at the 6<sup>th</sup> ACM International Conference on Web Search and Data Mining (WSDM 2013). Feb 2013.
- *Reliable Crowdsourced Preferences and Rankings*. Invited talk at NYU Courant Institute of Mathematical Sciences (Feb 2013).
- *Mining and Using Contextual Information from Large-Scale Web Search Logs*. Invited talk at Emory University (Nov 2012). Also presented at Georgia Institute of Technology.
- *Improving Classification in Large Taxonomies*. Invited talk at Amazon Machine Learning Lunch Seminar (May 2012). PDF version of slides are [available](#).

- *Preference, Consensus, and Choice in Crowdsourced Relevance*. Invited talk at UT-Austin, Information School, Crowdsourcing Seminar (March 2011). PDF version of slides are [available](#).
- *Class-Based Contextualized Search*. Invited talk at Carnegie Mellon, Language Technologies Institute Colloquium (Oct 2010). Also, presented at the University of Texas at Austin, Purdue University, and UIUC. PDF of recent slides are [available](#).
- *Leveraging Locality for Text Classifier Combination*  
IJCAI 2003 Doctoral Consortium
- *Using Asymmetric Distributions to Improve Classifier Probability Estimates*  
NIPS 2004 Workshop on Calibration and Probabilistic Prediction in Supervised Learning

## Research Experience

- September 2006 to Present  
Senior Researcher, *Context, Learning, and User Experience for Search (CLUES)*, Microsoft Research  
As a researcher, I have crafted an ambitious agenda in machine learning and information retrieval that has both a scientific and practical impact. I have been able to achieve such a balance – with more than 5 deployed technologies, 15 patents with 10 more patent applications filed and in process, 39 full papers in competitive venues (acceptance  $\leq$  25%), more than 25 additional peer reviewed papers in addition to directional non-peer reviewed papers (e.g. our diversity workshop summary has been influential in directing the field in separating types of diversity that require different techniques) while still providing external organizational leadership and remaining engaged in professional service to the community. My research agenda has been focused on the learning challenges posed by large-scale web and interaction data, and together with my collaborators, I have made several deep investigations in these areas. The majority of my work has focused on the applications of intelligent virtual assistants as well as contextual and personal search with the aims of identifying the primary signals that most aid in contextualization, modeling contextual signals and extracting representations that can be used at scale while preserving their uncertain nature, and generalizing to a framework that can be readily applied to both new signals and other domains. In addition to demonstrating how to use novel signals such as reading level and location of interest, our models have improved over state-of-the-art commercial baselines while demonstrating how previous disparate signals can be unified within one contextualization framework. In addition, I have developed text classification techniques based on machine learning insights that increase accuracy by 20-30% over previous state-of-the-art and demonstrated how text classification techniques can be used across both documents and queries to improve state-of-the-art performance of retrieval systems at scale (a goal chased for well more than a decade but never achieved over broad query sets seen in web retrieval systems). To handle the challenge of noisy gold truth, I have used human computation techniques to not only handle the lack of consensus in gold-standard data needed for learning but to consider techniques that progress toward explicit personalized judgments. Additionally, I have advanced research in virtual assistants by improving the context in which users interact with assistants – primarily focusing on time and location. This is related to efforts where I played a key role in shipping technology within Cortana to assist users by automatically identifying from e-mail commitments they made to others but are likely to forget – providing a very natural and human style of assistance. Finally, I have developed active learning techniques targeted at selecting both queries and documents for the ranking space, active evaluation techniques to handle challenges from accurately evaluating classifiers at scale where data in the wild deviates from comparatively small labeled sets, and active sampling techniques that can balance the real-time cost of acquiring a label with the value it delivers.
- June 2006 to August 2006  
Consulting Services for a health system and associated business partners  
Dr. Jaime Carbonell and I were contracted for consulting services to perform evaluations of medical NLP software for a large health system for the purpose of semi-automating the coding of patient records into ICD-9 codes (i.e., insurance and diagnosis codes). This coding process is typically a costly process that requires highly trained workers to spend significant time per record -- often a half hour or more for complex in-patient visit records. Undercoding leads to underbilling and loss of legitimately earned revenue; overcoding leads to overbilling and possible charges of fraud (with associated penalties). In addition, we also provided data-driven statistical machine learning methods in order to establish a baseline performance benchmark for systematic comparative quantitative evaluation of the medical NLP systems which passed initial examination.
- September 2005 to September 2006  
Chief Learning Architect, RADAR Project (Supervisor: Dr. Carbonell)  
My work on the RADAR project focused on statistically analyzing and explaining the benefits of learning from the year two evaluation of the RADAR system and advising on the direction of future learning efforts throughout the project. Additionally, I developed classification systems for automatically detecting both e-mails and sentences containing action-items or to-do items which require the user's attention. RADAR is a large research initiative whose goal is to develop a "cognitive assistant" centered on managing a user's tasks and priorities via e-mail, automatically generated briefings, room reservation, etc.



- September 1998 to May 2006  
Thesis Research  
With Drs. Jaime Carbonell and John Lafferty, I investigated methods to combine classifiers in a manner that makes use of their local reliability, variance, and dependence over the input space. In particular, my dissertation focused on developing a meta-classification scheme, STRIVE, which used the outputs of classifiers (decision trees, kNN, linear SVMs, and naive Bayes) in conjunction with a set of reliability-indicators I defined. The resulting model used the training data to automatically detect regions of poor classifier reliability and generate a more reliable combined prediction — outperforming standard approaches such as weighted combinations of the classifier outputs.
- Spring 2004 to July 2005  
Radar Project: Detecting Action-Items in E-mail  
With Dr. Jaime Carbonell, we defined a learning task aimed at reducing a user's time spent processing e-mail by detecting sentences that explicitly request information or action from a user. The primary work was performed in Spring - Summer 2004. I was responsible for adapting standard learning approaches that could effectively learn to detect not only e-mail containing action-items but particular sentences of interest within the e-mail. My responsibilities included: construction, design, and empirical evaluation of the machine learning methods that would serve as the prediction backbone; conducting labeled data collection and analysis in co-ordination with Jill Lehman; coordinating design of the User Interface and its interaction with the learning agent with Django Wexler and Jill; and generating balanced datasets as well as coordinating other aspects involving the learning agent for the HCI experiments conducted by Aaron Steinfeld and Jill. Continuing work involves publication of the findings, defining future directions (such as inductive transfer of learned models to reduce labeled data needs for new users), and performing ongoing experimental investigations.
- June 2002 to May 2003  
EBMT (Example-Based Machine Translation) Research Group  
Under the supervision of Dr. Jaime Carbonell, we investigated using overlap in candidate phrases output by an EBMT engine to guide the search for the optimal sentence construction from those phrases. I was responsible for developing efficient search techniques based on the heuristic function we used, coding and maintaining the prototype for overlap experimentation, and aiding in empirical analysis. The research indicated a consistent improvement in translation quality when using overlap to guide the search.
- August 2000 to December 2000 and May 2001 to August 2001  
Research Internship at Microsoft Research  
With Drs. Susan T. Dumais and Eric Horvitz, I investigated methods to combine classifiers using reliability predictions based upon context-sensitive indicators. We showed significant improvement over the whole space of linear utility functions for some domains while not negatively impacting other domains.
- February to May 1999  
Semester Project  
With Dr. Lafferty, developed a greedy-search incremental algorithm for multi-label classification tasks (i.e. where each instance can have more than one label) that can be applied to any classifier which outputs a posterior distribution.
- October to December 1998  
Semester Project  
With Dr. Roni Rosenfeld, attempted to find useful features to use in a maximum entropy language model. Focus turned towards a bias discovered in typical language modeling methods that cause the likelihood of underestimating a feature's probability to increase as the frequency of the feature increases.
- September 1997 to May 1998  
Undergraduate Honors Thesis  
With Dr. Ray Mooney, investigated applications of text categorization to recommender systems. The project focused on adaptations of Naive Bayes as well as evaluation methods for ranked data.
- April 1997 to August 1998  
Cycorp, Inc.  
Position: Programmer/Ontological Engineer
  - With Dr. Ken Murray, designing tests that helped characterize the subset of first-order logic that can be reliably transformed to a canonical form (including variable renaming and predicate reordering).
  - With Dr. Murray, developed semantics for certain 2<sup>nd</sup>-order relations as well as implementing structures & algorithms that allowed fast inference over these & similar specialized 2<sup>nd</sup>-order relations.
  - With Karen Pittman, defining semantics that allow one relation to be treated "virtually" as a different relation in a well-defined and cohesive way in order to display information appropriate to the user's technical level.
- January to May 1997  
Semester Project  
With Dr. Robert L. Causey and Dr. Ben Kuipers, experimented with constructing Bayesian Networks (via HUGIN and MSBN) to model real world phenomena. Culmination of the project was a paper summarizing an attempt to model a real-world court case (where much of the evidence was questionable, tainted, etc.) with a network.

## Teaching Experience

- Guest Lecturer, New York University, Stern School of Business, Data Mining For Business Analytics, March 5, 2013, Introduction to Text Classification and Applications.
- Guest Lecturer, University of North Carolina, Readings in Interactive Information Retrieval (INLS 696), October 4, 2012, Short- and Long-Term Personalization.
- ECIR 2010 Tutorial on [Machine Learning in IR: Recent Successes and New Opportunities](#)  
Paul Bennett, Misha Bilenko, Kevyn Collins-Thompson  
In *ECIR 2010*, Milton Keynes, UK, March, 2010.
- ICML 2009 Tutorial on [Machine Learning in IR](#)  
Paul Bennett, Misha Bilenko, Kevyn Collins-Thompson  
In *ICML 2009*, Montreal, Canada, June, 2009.
- Eberly Center for Teaching Excellence, Completion of Documentation of Teaching Development Program, Carnegie Mellon University, 2001-2006:
  - Participation in six seminars on central teaching activities:
    - Overview of Student Motivation
    - Conducting Productive Discussions
    - Creating Effective Assignments and Exams
    - Writing in the Disciplines
    - Reflecting on Your Teaching Style: Microteaching Workshop
    - Debunking Myths about Learning Styles
  - Observation and formal feedback on teaching on two occasions.
  - Design, preparation, and feedback on a syllabus for a course expected to teach (Machine Learning Graduate Level).
  - Design, presentation, and feedback of a special project -- Collaboration, Cheating, and Cultural Expectations.
- Guest Lecturer, University of Washington, Principles of Information Retrieval Systems (IMT 542), Introduction to Text Classification and Filtering, November 18, 2006.
- Guest Lecturer, Carnegie Mellon University, Tepper School of Business, Web-Based Information Architectures (MSEC 20-760), Introduction to Text Categorization, July 23, 2002.
- Teaching Assistant for PhD introductory course, "15-781: Machine Learning". This involved helping to determine course content, regular recitations, creating and preparing homework assignments and exams, and grading assignments.
- Teaching Assistant for "15-212: Fundamental Structures of Computer Science II". This involved teaching weekly recitations as well as helping to prepare and grade assignments.

## Other Academic Experience

- CMU Computer Science Department Graduate Admissions Committee, 2002-2003
- CMU CS Speaking Club, Spring 2001 - Fall 2005  
The Speaking Club is responsible for judging whether students have met the standards required to pass the speaking requirement. Two faculty and one student member attends each speaking requirement talk and a pass is required from each for the student to pass. The goal of the club is to help each student develop the skills required to present a well-organized, coherent talk to a general CS audience.

## Dissertation Committees

- Tobias Schnabel, Cornell University (Advisor: Thorsten Joachims). Computer Science Department. *Designing Learning Algorithms and Human Behavior for Interactive Systems*.
- Hyunjoon Jung, University of Texas at Austin (Advisor: Matt Lease). School of Information. *Temporal Modeling of Crowd Work Quality for Quality Assurance in Crowdsourcing*. Defended, November, 2015.
- Karthik Raman, Cornell University (Advisor: Thorsten Joachims). Computer Science Department. *Machine Learning from Human Preferences and Choices*. Defended, June, 2015.
- Joel Pfeiffer, Purdue University (Advisor: Jennifer Neville). Computer Science Department. *Overcoming Uncertainty for Within-Network Statistical Relational Learning*. Defended, December 2014.

## Professional Service

- Journal Editorial Service:
  - Guest Editor-in-Chief, ACM Transactions on Information Systems, Special Issue on Contextual Search and Recommendation (January 2015).
- Conference Organization:
  - Co-Chair SIGIR 2018 Short Papers
  - General Co-Chair [WSDM 2016](#).

- [TREC 2014 Web Track](#) Co-organizer.
- AAAI Conference on Human Computation and Crowdsourcing (HComp) Steering Committee (2013-2014).
- Co-organizer of the CIKM 2013 Workshop on [Exploiting Semantic Annotations for Information Retrieval](#) (ESAIR '13).
- [TREC 2013 Web Track](#) Co-organizer.
- Co-organizer of the ICML 2013 workshop on where [Machine Learning Meets Crowdsourcing](#).
- Research Program Co-chair of [CrowdConf 2013](#).
- Co-organizer of the CIKM 2013 Workshop on [Exploiting Semantic Annotations for Information Retrieval \(ESAIR '13\)](#).
- Co-organizer of the ICML 2013 Workshop on [Machine Learning Meets Crowdsourcing](#).
- [CIKM 2012](#) Co-Demo Chair.
- [WSDM 2012](#) Tutorial Chair.
- Co-organizer of the SIGIR 2011 Workshop on [Enriching Information Retrieval](#).
- Co-organizer of the SIGIR 2009 Workshop on [Redundancy, Diversity, and Interdependent Document Relevance](#).
- Co-organizer of the KDD 2009 Workshop on [Human Computation](#).
- Co-organizer and Chair of the [SIGIR 2008 Workshop on Beyond Binary Relevance: Preferences, Diversity, and Set-Level Judgments](#).
- Publication Chair of ICML, 2006.
- Senior Program Committees:
  - CIKM SPC 2011-2013.
  - HCOMP 2013 Best Paper Committee.
  - SIGIR Area Chair, 2013-2015, 2017.
  - ICML Area Chair, 2015.
  - SIGKDD SPC 2014-2017.
  - WSDM 2018 SPC
- Conference Program Committees:
  - AAAI Doctoral Consortium 2008-2010.
  - ACL 2008.
  - CHI 2009.
  - CIKM 2010.
  - ECIR 2008-2010.
  - ECML 2008 Workshop on Preference Learning.
  - ECML 2009-2010.
  - EMNLP 2012.
  - HLT-NAACL 2007.
  - HCOMP 2009-2011, 2013.
  - ICML 2010-2013.
  - IIX 2012.
  - SDM 2013.
  - SIGIR 2004-2011.
  - SIGIR 2011 Crowdsourcing for IR Workshop
  - SIGKDD 2011, 2013.
  - WSDM 2011, 2013, 2015.
  - WWW 2010-2013.
- Journal Reviewing
  - Reviewer for the Journal of Artificial Intelligence Research (JAIR).
  - Reviewer for the Journal of the American Society for Information Science and Technology (JASIST).
  - Reviewer for the Journal of Machine Learning Research (JMLR).
  - Reviewer for ACM Transactions on Information Systems (TOIS).
  - Reviewer for ACM Transactions on Internet Technology (TOIT).
  - Reviewer ACM Transactions on the Web (TWeb).
- Co-organizer [IR Seminar Series](#), Summer 2004 - Summer 2005.

## Other Employment

Other work experience information (including tutoring, technical support, acting, directing, scenic design, ....) from 1992-1997 available upon request.

## Awards & Honors

- ACM UMAP 2016 Best Student Paper Award
- WSDM 2015 Outstanding Reviewer Award

- SIGIR 2013 Best Student Paper Award
- SIGIR 2012 Best Paper Honorable Mention (Only 3 paper awards of any type given)
- National Science Foundation Graduate Fellowship, Honorable Mention (1999)
- University Honors (Spring 1995 - Spring 1997)
- Election to Phi Beta Kappa Honor Society
- Merit Scholarships held during Undergraduate Study
  - Dedman Merit Scholarship
  - National Merit Scholarship
  - Nelda C. and H.J. Lutcher Stark Foundation Scholarship
  - Small Schools Academic Challenge Scholarship
  - V.F.W. Voice of Democracy Scholarship
  - Wimberley Lion's Club Scholarship
- Completion of Computer Science Honors Program
- Completion of Plan II Honors Program

## Patents

- *US Pat. No. 9594837. Prediction and Information Retrieval for Intrinsically Diverse Sessions*  
Inventors: Paul N. Bennett, Kevyn B. Collins-Thompson, Paul N. Bennett  
Issue Date: March 2017.
- *US Pat. No. 9535995. Optimizing a Ranker for a Risk-Oriented Objective*  
Inventors: Kevyn B. Collins-Thompson, Paul N. Bennett, Lidan Wang  
Issue Date: January, 2017.
- *US Pat. No. 9064016, Ranking Search Results Using Result Repetition*  
Inventors: Milad Shokouhi, Ryen W. White, Paul N. Bennett  
Issue Date: June, 2015.
- *US Pat. No. 9026367, Dynamic Destination Navigation System*  
Inventors: Paul N. Bennett, Tim S. Paek, Victor Bahl  
Issue Date: May, 2015.
- *US Pat. No. 9020936, Using Categorical Metadata to Rank Search Results*  
Inventors: Susan T. Dumais, Paul N. Bennett, and Krysta M. Svore  
Issue Date: April, 2015.
- *US Pat. No. 8972399, Ranking Based On Social Activity Data*  
Inventors: Paul N. Bennett, Emre M. Kiciman, Peter R Bailey, Huizhong Duan, Nikhil Dandekar  
Issue Date: March, 2015.
- *US Pat. No. 8954423, Using Reading Levels in Responding to Requests*  
Inventors: Kevyn B Collins-Thompson, Paul N. Bennett, Ryen W. White, Sebastian de la Chica, David A. Sontag  
Issue Date: Feb. 2015.
- *US Pat. No. 8700544, Functionality for Personalizing Results*  
Inventors: Kevyn B Collins-Thompson, Susan T. Dumais, Paul N. Bennett, Ryen W. White, David A. Sontag  
Issue Date: April 2014.
- *US Pat. No. 8825641, Measuring Duplication in Search Results*  
Inventors: Paul N. Bennett, Filip Radlinski, and Emine Yilmaz  
Issue Date: September, 2014.
- *US Pat. No. 8719249, Query Classification*  
Inventors: Kevyn B Collins-Thompson, Max M. Chickering, Susan T. Dumais, Paul N. Bennett, and Dan J Liebling  
Issue Date: May, 2014.
- *US Pat. No. 8645289, Structured Cross-lingual Relevance Feedback For Enhancing Search Results*  
Inventors: Jianfeng Gao, Paul N. Bennett, Jagadeesh Jagarlamudi, and Kristen Parton  
Issue Date: February, 2014.
- *US Pat. No. 8396875, Online Stratified Sampling for Classifier Evaluation*  
Inventors: Paul N. Bennett and Vitor R. Carvalho  
Issue Date: March, 2013
- *US Pat. No. 8296330, Hierarchical Classification*  
Inventors: Paul N. Bennett and Nam H Nguyen  
Issue Date: October, 2012.
- *US Pat. No. 8069179, Preference Judgments For Relevance*  
Inventors: Max M. Chickering, Susan T. Dumais, Paul N. Bennett, and Ben Carterette  
Issue Date: November, 2011.

- *US Pat. No. 7107254, Probabilistic Models and Methods for Combining Multiple Content Classifiers*  
Inventors: Susan T. Dumais, Eric J. Horvitz, and Paul N. Bennett  
Issue Date: September, 2006.

## Patent Applications in Process

### Patents Published

- *Ranking Content Items Based on Preference Scores*  
Inventors: Paul N. Bennett, Milad Shokouhi, Richard A. Caruana  
Filed: May, 2015. Published: Nov, 2016. Publication number: US-2016-0342692-A1.
- *Automatic Extraction of Commitments and Requests from Communications and Content*  
Inventors: Paul N. Bennett, Nikrouz Ghotbi, Eric J. Horvitz, Richard L. Hughes, Pradeep Singh, Ryen W. White.  
Filed: May, 2015. Published: Nov, 2016. Publication number: US-2016-0335572-A1.
- *Management of Commitments and Requests Extracted from Communications and Content*  
Inventors: Paul N. Bennett, Nirupama Chandrasekaran, Michael Gamon, Nikrouz Ghotbi, Eric J. Horvitz, Richard L. Hughes, Pradeep Singh, Ryen W. White.  
Filed: May, 2015. Published: Nov, 2016. Publication number: US-2016-0337295-A1.
- *Modeling Actions, Consequences, and Goal Achievement from Social Media and Other Digital Traces*  
Inventors: Emre Kiciman, Paul N. Bennett, Jaime B. Teevan, Susan T. Dumais.  
Filed: Dec, 2014. Published: Jun, 2016. Publication number: US-2016-0171063-A1.
- *Learning and Using Contextual Content Retrieval Rules for Query Disambiguation*  
Inventors: Paul N. Bennett, Kevyn Collins-Thompson, Siranush Sarkizova, Milad Shokouhi, Marc Sloan  
Filed: June, 2014. Published: Dec, 2015. Publication number: US-2015-0363485-A1
- *Using Aggregate Location Metadata to Provide a Personalized Service*  
Inventors: Paul N. Bennett, Filip Radlinski, Ryen W. White, Emine Yilmaz  
Filed: June, 2011. Published: December, 2012. Publication number: US-2012-0317104-A1
- *Modeling Intent and Ranking Search Results Using Activity-based Context*  
Inventors: Susan T. Dumais, Paul N. Bennett, Xiaoyuan Cui, Ryen W. White, Peter R Bailey, and Fedor V Borisjuk  
Filed: December, 2010. Published: June, 2012. Publication number: US-2012-0158685-A1.

### Patents Filed

Patents are *published* to the public 18 months after the date filed. Until then, the application is confidential. Inventor on 3 patent applications filed but not yet published. File dates:

- December, 2016
- July, 2016
- October, 2015

## References

Available upon request.