

Panos Ipeirotis

PERSONAL DATA

Information Systems Group
Department of Information, Operations and Management Sciences
Leonard N. Stern School of Business
New York University
44 West 4th Street, Suite 8-84
New York, NY 10012-1126, USA

Phone: +1 (212) 998-0803
Fax: +1 (212) 995-4228
Email: panos@stern.nyu.edu
Website: <http://www.ipeirotis.com>

EDUCATION

- Sep 2004: Ph.D. (with distinction), Computer Science, Columbia University
- Feb 2003: M.Phil., Computer Science, Columbia University
- Feb 2001: M.Sc., Computer Science, Columbia University
- Jul 1999: Diploma, Computer Engineering and Informatics, University of Patras, Greece

PROFESSIONAL EMPLOYMENT

- Jan 2016 - now: Professor of Information Systems, Department of Information, Operations, and Management Sciences, Leonard N. Stern School of Business, New York University, New York, NY, USA
- Dec 2015 - now: Professor of Data Science, Center for Data Science, New York University, New York, NY, USA
- Oct 2005 - now: Affiliated Faculty (Courtesy Appointment), Computer Science Department, Courant Institute of Mathematical Sciences, New York University, New York, NY, USA
- May 2014 - May 2015: Consultant, World Bank, Washington D.C., United States
- Feb 2013 - Sep 2014: Visiting Scientist, Google, New York, NY, USA and Mountain View, CA, USA
- Jan 2012 - Dec 2012: Academic-in-residence, oDesk (now UpWork), Redwood City, CA, USA
- Jun 2011 - Feb 2013: Chief Scientist, Tagasauris, New York, NY, USA
- Sep 2010 - Dec 2015: Associate Professor, Department of Information, Operations, and Management Sciences, Leonard N. Stern School of Business, New York University, New York, NY, USA
- Sep 2004 - Aug 2010: Assistant Professor, Department of Information, Operations, and Management Sciences, Leonard N. Stern School of Business, New York University, New York, NY, USA
- May 2003 - Aug 2003: Research Intern, Microsoft Research Ltd., Cambridge, U.K
- Sep 1999 - Aug 2004: Graduate Research Assistant, Computer Science Department, Columbia University, New York, NY, USA
- Sep 1998 - Jun 1999: Research Assistant, Graphics and Multimedia Lab, Computer Engineering Informatics Department, University of Patras, Patras, Greece

- Nov 1997 - Aug 1999: Software Engineer, Computing and Networking Services, Computer Technology Institute, Patras, Greece
- Oct 1995 - Nov 1997: System Administrator, Computer Engineering and Informatics Department, University of Patras, Patras, Greece
- Sep 1992 - Aug 1994: Instructor, Data Computers, Serres, Greece

AWARDS AND GRANTS

- Jun 2016: Outstanding Reviewer Award; Twenty-fifth International Joint Conference on Artificial Intelligence (IJCAI-16)
- Apr 2016: Senior Member Promotion; IEEE
- Oct 2015: Lagrange Prize; CRT Foundation
- Jan 2015: Outstanding Reviewer Award; Eighth ACM International WSDM Conference (WSDM 2015)
- Nov 2014: Best Paper for Years 2011-14; Management Science, ISS/INFORMS
- Nov 2014: Best Paper Award; Second AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2014)
- Nov 2014: Best Paper Award Runner-up; INFORMS 2014 Data Mining Student Competition
- Dec 2012: Google Focused Research Award, Never Failing Goggles, \$1,500,000, with Serge Belongie and Pietro Perona Google
- Dec 2012: Best Theme Paper Award; International Conference on Information Systems (ICIS 2012)
- Nov 2011: Google Research Award, Designing Ranking Systems for Products on Search Engines Using Economic Theory Models for Improving Online Shopping, \$60,000, with Anindya Ghose Google
- Apr 2011: Best Paper Award; 20th International World Wide Web Conference
- Jul 2010: Google Research Award, Collaborative Adaptive Data Sharing, \$70,000, with Vagelis Hristidis Google
- May 2010: NYU Research Challenge Grant, Comprehensive and High Quality Monitoring of the Web for Academic Citations, \$18,000 New York University
- Apr 2010: NYU-Poly Seed Grant, Understanding the structure of evolving textual collections, \$30,000, with Torsten Suel NYU-Poly
- Mar 2009: Marketing Science Institute Grant: , The economic impact of user-generated content: Combining text mining with demand estimation in the hotel industry, \$6,500, with Anindya Ghose Marketing Science Institute
- Aug 2008: Best Paper Award Runner Up; ACM International Conference on Knowledge Management and Data Mining (KDD 2008)
- Jun 2007: NET Institute Summer Grant, Deriving the Pricing Power of Product Features by Mining Consumer Reviews, \$4,500, with Anindya Ghose NET Institute
- Mar 2007: Microsoft Virtual Earth Award, Local Search for Hotel and Restaurants Using Econometrics Spatial Data and Image Classification, \$35,000, with Anindya Ghose Microsoft Research
- Jan 2007: CAREER Award, Towards a Text-Centric Database Management System, IIS-0643846, \$499,957 National Science Foundation
- Jun 2006: Best Paper Award; ACM International Conference on Management of Data (SIGMOD 2006)

- Jan 2006: Microsoft Live Labs Search Award, Combining Econometric and Text Mining Approaches for Measuring the Effect of Online Information Exchange, \$37,000, with Anindya Ghose Microsoft Live Labs
- Aug 2005: Departmental Nominee, ACM Doctoral Dissertation Award; Computer Science Department, Columbia University
- Apr 2005: Best Paper Award; IEEE International Conference on Data Engineering (ICDE 2005)
- Oct 2004: Award for Excellence in Teaching; School of Engineering, Columbia University
- Feb 2002: Graduate Fellowship; DIMACS (Center for Discrete Mathematics and Theoretical Computer Science)
- Jun 2001: Best Paper Award Finalist; ACM+IEEE Joint Conference on Digital Libraries (JCDL 2001)
- May 2001: Ph.D. Teaching Award of Excellence; Computer Science Department, Columbia University
- Jan 2001: Honorary Distinction Award; Technical Chamber of Greece
- Sep 2000: Award for Excellence in Teaching; School of Engineering, Columbia University
- Jan 2000: Honorary Distinction Award; Technical Chamber of Greece
- Sep 1999: Fellowship; , \$10,000 Empeirikio Foundation
- May 1999: Five Fellowships from Greek Scholarship Foundation;
- Jan 1994: Member of the Greek Olympiad team; 26th International Chemistry Olympiad (Oslo 1994)
- Jan 1994: Gold Medal; 8th National Greek Competition in Chemistry

TEACHING EXPERIENCE

- Spring 2018: Projects in Programming and Data Science (INFO-UB.0024.01); Instructor, Enrollment: 40; IOMS Department, Stern School of Business, New York University, New York
- Fall 2017: Projects in Programming and Data Science (INFO-UB.0024.01); Instructor, Rating: 6.3/7.0, Enrollment: 40; IOMS Department, Stern School of Business, New York University, New York
- Fall 2017: Introduction to Programming and Data Science (INFO-UB.0023.01 and INFO-UB.0023.02); Instructor, Rating: 6.4/7.0, Enrollment: 160; IOMS Department, Stern School of Business, New York University, New York
- Summer 2017: Dealing with Data (XBA1-GB.8217.84); Instructor, Rating: 6.2/7.0, Enrollment: 70; Business Analytics Program, Stern School of Business, New York University, New York
- Spring 2017: Projects in Programming and Data Science (INFO-UB.0024.01); Instructor, Rating: 6.8/7.0, Enrollment: 45; IOMS Department, Stern School of Business, New York University, New York
- Fall 2016: Introduction to Programming and Data Science (INFO-UB.0023.01, INFO-UB.0023.02, and INFO-UB.0023.03); Instructor, Rating: 6.5/7.0, Enrollment: 148; IOMS Department, Stern School of Business, New York University, New York
- Summer 2016: Dealing with Data (XBA1-GB.8217.84); Instructor, Rating: 5.4/7.0, Enrollment: 70; Business Analytics Program, Stern School of Business, New York University, New York
- Spring 2016: Projects in Programming and Data Science (INFO-UB.0024.01); Instructor, Rating: 6.3/7.0, Enrollment: 27; IOMS Department, Stern School of Business, New York University, New York

- Winter 2016: Dealing with Data (INFO-GB.2346.W1); Instructor, Rating: 6.0/7.0, Enrollment: 40; IOMS Department, Stern School of Business, New York University, New York
- Fall 2015: Introduction to Programming and Data Science (INFO-UB.0023.01); Instructor, Rating: 6.3/7.0, Enrollment: 38; IOMS Department, Stern School of Business, New York University, New York
- Summer 2015: Dealing with Data (XBA1-GB.8217.84); Instructor, Rating: 6.1/7.0, Enrollment: 60; Business Analytics Program, Stern School of Business, New York University, New York
- Winter 2015: Dealing with Data (INFO-GB.2346.0W1); Instructor, Rating: 6.5/7.0, Enrollment: 27; IOMS Department, Stern School of Business, New York University, New York
- Fall 2014: Dealing with Data (INFO-GB.2346.10); Instructor, Rating: 6.0/7.0, Enrollment: 41; IOMS Department, Stern School of Business, New York University, New York
- Fall 2014: Dealing with Data (INFO-UB.0046.01); Instructor, Rating: 7.0/7.0, Enrollment: 11; IOMS Department, Stern School of Business, New York University, New York
- Fall 2011: Information Technology in Business and Society (INFO-UB.0001.004, INFO-UB.0001.005, and INFO-UB.0001.006); Instructor, Rating: 6.2/7.0, Enrollment: 158; IOMS Department, Stern School of Business, New York University, New York
- Fall 2010: Information Technology in Business and Society (C20.0001.002 and C20.0001.003); Instructor, Rating: 5.4/7.0, Enrollment: 94; IOMS Department, Stern School of Business, New York University, New York
- Fall 2009: Information Technology in Business and Society (C20.0001.004 and C20.0001.005); Instructor, Rating: 6.3/7.0, Enrollment: 106; IOMS Department, Stern School of Business, New York University, New York
- Fall 2008: Search and the New Economy (B20.3136.010); Instructor, Rating: 5.7/7.0, Enrollment: 69; IOMS Department, Stern School of Business, New York University, New York
- Fall 2008: Information Technology in Business and Society (C20.0001.002 and C20.0001.003); Instructor, Rating: 6.4/7.0, Enrollment: 109; IOMS Department, Stern School of Business, New York University, New York
- Winter 2008: Search and the New Economy (B20.3136.0W1); Instructor, Rating: 6.4/7.0, Enrollment: 37; IOMS Department, Stern School of Business, New York University, New York
- Fall 2007: Information Technology in Business and Society (C20.0001.002, C20.0001.005); Instructor, Rating: 6.5/7.0, Enrollment: 100; IOMS Department, Stern School of Business, New York University, New York
- Fall 2007: Information Systems Technical Research Seminar (B20.3383.020); Instructor, Rating: 7.0/7.0, Enrollment: 1; IOMS Department, Stern School of Business, New York University, New York
- Fall 2006: Information Technology in Business and Society (C20.0001.003 and C20.0001.004); Instructor, Rating: 6.6/7.0, Enrollment: 80; IOMS Department, Stern School of Business, New York University, New York
- Spring 2006: Information Systems Technical Research Seminar (B20.3383.020); Instructor, Rating: 7.0/7.0, Enrollment: 4; IOMS Department, Stern School of Business, New York University, New York
- Spring 2006: Information Technology in Business and Society (C20.0001.001 and C20.0001.002); Instructor, Rating: 5.7/7.0, Enrollment: 80; IOMS Department, Stern School of Business, New York University, New York
- Spring 2005: Computer-based Systems for Management Support (C20.0001.003 and C20.0001.004); Instructor, Rating: 5.6/7.0, Enrollment: 112; IOMS Department, Stern School of Business, New York University, New York

- Spring 2003: Advanced Database Systems (COMS 6111); Teaching Assistant; Computer Science Department, Columbia University, New York
- Spring 2002: Advanced Database Systems (COMS 6111); Instructor; IOMS Department, Stern School of Business, New York University, New York
- Spring 2001: Advanced Database Systems (COMS 6111); Teaching Assistant; Computer Science Department, Columbia University, New York
- Spring 2001: Advanced Database Systems (COMS 6111); Instructor; Columbia Video Network, Computer Science Department, Columbia University, New York
- Summer 2000: Advanced Database Systems (COMS 6111); Instructor; Columbia Video Network, Computer Science Department, Columbia University, New York
- Spring 2000: Advanced Database Systems (COMS 6111); Teaching Assistant; Computer Science Department, Columbia University, New York
- Spring 1996: Programming Languages (); Teaching Assistant; Computer Engineering and Informatics Department, University of Patras, Greece
- Fall 1992: Programming Languages and Operating Systems (); Instructor; Data Computer School, Serres, Greece

PUBLICATIONS

Journals

- **Jan 2018: Modeling Consumer Footprints on Search Engines: An Interplay with Social Media;** Beibei Li, Anindya Ghose, and Panagiotis Ipeirotis; Management Science, Forthcoming. Citations: 16
- **Mar 2017: Cost-Effective Quality Assurance in Crowd Labeling;** Jing Wang, Panagiotis G. Ipeirotis, and Foster Provost; Information Systems Research, Volume 28, Number 1, March 2017. Citations: 28
- **Jun 2016: Reputation Transferability in Online Labor Markets;** Marios Kokkodis and Panagiotis Ipeirotis; Management Science, Volume 62, Issue 6, Pages 1687-1706. Citations: 14
- **Oct 2015: A System for Scalable and Reliable Technical-Skill Testing in Online Labor Markets;** Maria Christoforaki and Panos Ipeirotis; Computer Networks, Volume 90, Pages 110-120. Citations: 2
- **Feb 2015: Beat the Machine: Challenging Humans to Find a Predictive Models Unknown Unknowns;** Joshua Attenberg, Panagiotis G. Ipeirotis, and Foster Provost; ACM Journal on Data and Information Quality, Volume 6, Issue 1, March 2015. Citations: 37
- **Mar 2014: Examining the Impact of Ranking on Consumer Behavior and Search Engine Revenue;** Anindya Ghose, Panagiotis G. Ipeirotis, and Beibei Li; Management Science, Volume 60, Issue 7, July 2014. Citations: 54
- **Mar 2014: Repeated Labeling Using Multiple Noisy Labelers;** Panagiotis G. Ipeirotis, Foster Provost, Victor S. Sheng, and Jing Wang; Data Mining and Knowledge Discovery, March 2014, Volume 28, Issue 2, pp 402-441. Citations: 75
- **Feb 2014: Facilitating Document Annotation using Content and Querying Value;** Eduardo J. Ruiz, Vagelis Hristidis, and Panagiotis G. Ipeirotis; IEEE Transactions on Knowledge and Data Engineering (TKDE). Citations: 23
- **May 2012: Designing Ranking Systems for Hotels on Travel Search Engines by Mining User-Generated and Crowd-Sourced Content;** Anindya Ghose, Panagiotis G. Ipeirotis, and Beibei Li; Marketing Science, Vol. 31, No. 3, May/June 2012. Citations: 269

- **Feb 2012: Answering General Time-Sensitive Queries;** Wisam Dakka, Luis Gravano, and Panagiotis G. Ipeirotis; IEEE Transactions on Knowledge and Data Engineering (TKDE), Vol. 24, No. 2, February 2012. Citations: 135
- **Oct 2011: Relevance-Based Retrieval on Hidden-Web Text Databases without Ranking Support;** Vagelis Hristidis, Yuheng Hu, and Panagiotis G. Ipeirotis; IEEE Transactions on Knowledge and Data Engineering (TKDE), Vol. 23, No. 10, October 2011. Citations: 11
- **Oct 2011: Estimating the Helpfulness and Economic Impact of Product Reviews: Mining Text and Reviewer Characteristics;** Anindya Ghose and Panagiotis G. Ipeirotis; IEEE Transactions on Knowledge and Data Engineering (TKDE), Vol. 23, No. 10, October 2011. Citations: 542
- **Aug 2011: Deriving the Pricing Power of Product Features by Mining Consumer Reviews;** Nikolay Archak, Anindya Ghose, and Panagiotis G. Ipeirotis; Management Science, Vol. 57, No 8, Pages 1485-1509. Citations: 335
- **Dec 2010: Analyzing the Amazon Mechanical Turk Marketplace;** Panagiotis G. Ipeirotis; ACM XRDS (Crossroads), Vol. 17, No 2, Winter 2010. Citations: 424
- **Aug 2010: Running Experiments on Amazon Mechanical Turk;** Gabriele Paolacci, Jesse Chandler, and Panagiotis G. Ipeirotis; Judgment and Decision Making, Vol. 5, No. 5, August 2010. Citations: 2104
- **Apr 2009: A Quality-Aware Optimizer for Information Extraction;** Alpa Jain and Panagiotis G. Ipeirotis; ACM Transactions on Database Systems (TODS), Vol. 34, No. 1, April 2009. Citations: 38
- **Mar 2009: The EconoMining Project at NYU: Studying the Economic Value of User-Generated Content on the Internet;** Anindya Ghose and Panagiotis G. Ipeirotis; Journal of Revenue and Pricing Management, Vol. 8, No 2-3, March 2009. Citations: 40
- **Dec 2008: Building Query Optimizers for Information Extraction: The SQoUT Project;** Alpa Jain, Panagiotis G. Ipeirotis, and Luis Gravano; SIGMOD Record, Special Issue on Managing Information Extraction, Vol. 37, No. 4, December 2008. Citations: 16
- **Mar 2008: Classification-Aware Hidden-Web Text Database Selection;** Panagiotis G. Ipeirotis and Luis Gravano; ACM Transactions on Information Systems (TOIS), Vol. 26, No. 2, March 2008. Citations: 18
- **Nov 2007: Towards a Query Optimizer for Text-Centric Tasks;** Panagiotis G. Ipeirotis, Eugene Agichtein, Pranay Jain, and Luis Gravano; ACM Transactions on Database Systems (TODS), Vol. 32, No. 4, November 2007. Citations: 35
- **Aug 2007: Modeling and Managing Changes in Text Databases;** Panagiotis G. Ipeirotis, Alexandros Ntoulas, Junghoo Cho, and Luis Gravano; ACM Transactions on Database Systems (TODS), Vol. 32, No. 3, August 2007. Citations: 13
- **Jan 2007: Duplicate Record Detection: A Survey;** Ahmed K. Elmagarmid, Panagiotis G. Ipeirotis, and Vassilios S. Verykios; IEEE Transactions on Knowledge and Data Engineering (TKDE), Vol. 19, No. 1, January 2007. Citations: 1579
- **Jan 2003: QProber: A System for Automatic Classification of Hidden-Web Databases;** Luis Gravano, Panagiotis G. Ipeirotis, and Mehran Sahami; ACM Transactions on Information Systems (TOIS), Vol. 21, No. 1, January 2003. Citations: 155
- **Mar 2002: Query- vs. Crawling-Based Classification of Searchable Web Databases;** Luis Gravano, Panagiotis G. Ipeirotis, and Mehran Sahami; IEEE Data Engineering Bulletin, Vol. 25, No. 1, March 2002. Citations: 12
- **Dec 2001: Using q-grams in a DBMS for Approximate String Processing;** Luis Gravano, Panagiotis G. Ipeirotis, Hosagrahar Visvesvaraya Jagadish, Nick Koudas, Shanmugaelayut

Muthukrishnan, Lauri Pietarinen, and Divesh Srivastava; IEEE Data Engineering Bulletin, Vol. 24, No. 4, December 2001. Citations: 150

Conferences

- **Feb 2018: Demographics and Dynamics of Mechanical Turk Workers;** Djellel Difallah, Elena Filatova, and Panos Ipeirotis; WSDM 2018: The Eleventh ACM International Conference on Web Search and Data Mining; 83/514 = 16%
- **Dec 2016: Ridesharing and the Use of Public Transportation;** ; International Conference on Information Systems - Digital Innovation at the Crossroads, (ICIS 2016)
- **Jun 2015: Building a Bird Recognition App and Large Scale Dataset With Citizen Scientists: The Fine Print in Fine-Grained Dataset Collection;** Grant Van Horn, Steve Branson, Ryan Farrell, Scott Haber, Jessie Barry, Panos Ipeirotis, Pietro Perona, and Serge Belongie; The Twenty-Eighth IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2015). Citations: 15
- **May 2015: Getting More for Less: Optimized Crowdsourcing with Dynamic Tasks and Goals;** Ari Kobren, Chun How Tan, Panos Ipeirotis, and Evgeniy Gabrilovich; 24th International World Wide Web Conference (WWW 2015); 131/929 = 14.1% accepted. Citations: 5
- **May 2015: The Dynamics of Micro-Task Crowdsourcing The Case of Amazon Mechanical Turk;** Djellel Eddine Difallah, Michele Catasta, Gianluca Demartini, Panos Ipeirotis, and Philippe Cudré-Mauroux; 24th International World Wide Web Conference (WWW 2015); 131/929 = 14.1% accepted. Citations: 26
- **Feb 2015: Hiring Behavior Models for Online Labor Markets;** Marios Kokkodis, Panagiotis Papadimitriou, and Panagiotis G. Ipeirotis; Eighth ACM International Conference on Web Search and Data Mining (WSDM 2015); 40/238 = 16.8% accepted. Citations: 9
- **Dec 2014: The Utility of Skills in Online Labor Markets;** Marios Kokkodis and Panagiotis Ipeirotis; Thirty-Fifth International Conference on Information Systems (ICIS 2014). Citations: 5
- **Nov 2014: STEP: A Scalable Testing and Evaluation Platform;** Maria Christoforaki and Panos Ipeirotis; Second AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2014); 26/82 = 30% acceptance rate; Best Paper Award. Citations: 6
- **Jun 2014: Efficient Filtering on Hidden Document Streams;** Eduardo Ruiz, Vagelis Hristidis, and Panagiotis G. Ipeirotis; 8th International AAAI Conference on Weblogs and Social Media (ICWSM 2014); 64/279 = 23% accepted. Citations: 6
- **Jun 2014: Composing and Analyzing Crowdsourcing Workflows;** Panagiotis G. Ipeirotis, Greg Little, and Thomas W. Malone; Collective Intelligence 2014
- **Apr 2014: Quiz: Targeted Crowdsourcing with a Billion (Potential) Users;** Panos Ipeirotis and Evgeniy Gabrilovich; 23rd International World Wide Web Conference (WWW 2014); 84/650=12.9% accepted. Citations: 44
- **Feb 2014: Trust, but Verify: Predicting Contribution Quality for Knowledge Base Construction and Curation;** Chun How Tan, Eugene Agichtein, Evgeniy Gabrilovich, and Panos Ipeirotis; Seventh ACM Conference on Web Search and Data Mining (WSDM 2014); 64/356 = 18% accepted. Citations: 12
- **Feb 2013: Have you Done Anything Like That? Predicting Performance Using Inter-category Reputation;** Marios Kokkodis and Panagiotis G. Ipeirotis; Sixth International Conference on Web Search and Data Mining (WSDM 2013); 73/387=19% accepted. Citations: 19
- **Dec 2012: Bonus, Disclosure, and Choice: What Motivates the Creation of High-Quality Paid Reviews?;** Jing Wang, Anindya Ghose, and Panagiotis G., Ipeirotis; Thirty Third International Conference on Information Systems (ICIS 2012). Citations: 23

- **Dec 2012: Search Less, Find More? Examining Limited Consumer Search with Social Media and Product Search Engines;** Anindya Ghose, Panagiotis G. Ipeirotis, and Beibei Li; Thirty Third International Conference on Information Systems (ICIS 2012); Best Theme Paper Award. Citations: 5
- **Dec 2011: Content and Context: Identifying the Impact of Qualitative Information on Consumer Choice;** Sinan Aral, Panagiotis G. Ipeirotis, and Sean Taylor; Proceedings of the International Conference on Information Systems (ICIS), 2011. Citations: 18
- **Mar 2011: Towards a Theory Model for Product Search;** Beibei Li, Anindya Ghose, and Panagiotis G. Ipeirotis; Proceedings of the 20th International World-Wide Web Conference (WWW), 2011; 81/658=12% accepted; Best Paper Award. Citations: 33
- **Mar 2010: Ranked Queries Over Sources with Boolean Query Interfaces without Ranking Support;** Vagelis Hristidis, Yuheng Hu, and Panagiotis G. Ipeirotis; Proceedings of the 25th IEEE International Conference on Data Engineering (ICDE), 2010; 20% accepted. Citations: 8
- **Jan 2010: Designing Ranking Systems for Hotels on Travel Search Engines to Enhance User Experience;** Anindya Ghose, Panagiotis G. Ipeirotis, and Beibei Li; Proceedings of the International Conference on Information Systems (ICIS), 2010
- **Jul 2009: Modeling Volatility in Prediction Markets;** Nikolay Archak and Panagiotis G. Ipeirotis; Proceedings for the 10th ACM Conference on Electronic Commerce (EC), 2009; 40/158 = 25% accepted. Citations: 3
- **Mar 2009: Join Optimization of Information Extraction Output: Quality Matters;** Alpa Jain, Panagiotis G. Ipeirotis, An-Hai Doan, and Luis Gravano; Proceedings of the 25th IEEE International Conference on Data Engineering (ICDE), 2009 ; 93/554 = 16.7% accepted. Citations: 27
- **Feb 2009: Query by Document;** Yin Yang, Nilesch Bansal, Wisam Dakka, Panagiotis G. Ipeirotis, Nick Koudas, and Dimitris Papadias; Proceedings of the Second ACM International Conference on Web Search and Data Mining (WSDM), 2009 ; 29/170 = 17% accepted. Citations: 90
- **Aug 2008: Get Another Label? Improving Data Quality and Data Mining Using Multiple, Noisy Labelers;** Victor S. Sheng, Foster Provost, and Panagiotis G. Ipeirotis; Proceedings of the Fourteenth ACM International Conference on Knowledge Discovery and Data Mining (KDD), 2008 ; 50/ 500 < 10% accepted; Best Paper Award Runner Up. Citations: 672
- **Apr 2008: Automatic Extraction of Useful Facet Hierarchies from Text Databases;** Wisam Dakka and Panagiotis G. Ipeirotis; Proceedings of the 24th IEEE International Conference on Data Engineering (ICDE), 2008 ; 75/617 = 12.1% accepted. Citations: 106
- **Aug 2007: Designing Novel Review Ranking Systems: Predicting Usefulness and Impact of Reviews;** Anindya Ghose and Panagiotis G. Ipeirotis; Proceedings of the Ninth International Conference on Electronic Commerce (ICEC), 2007. Citations: 272
- **Aug 2007: Show me the Money Deriving the Pricing Power of Product Features by Mining Consumer Reviews;** Nikolay Archak, Anindya Ghose, and Panagiotis G. Ipeirotis; Proceedings of the Thirteenth ACM International Conference on Knowledge Discovery and Data Mining (KDD), 2007 ; 100/513 < 20% accepted. Citations: 205
- **Jun 2007: Opinion Mining Using Econometrics: A Case Study on Reputation Systems;** Anindya Ghose, Panagiotis G. Ipeirotis, and Arun Sundararajan; Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics (ACL), 2007 ; 132/588 = 22% accepted. Citations: 142
- **Jun 2006: To Search or to Crawl? Towards a Query Optimizer for Text-Centric Tasks;** Panagiotis G. Ipeirotis, Eugene Agichtein, Pranay Jain, and Luis Gravano; Proceedings of the 2006 ACM International Conference on Management of Data (SIGMOD), 2006; 58/446=13% accepted; Best Paper Award. Citations: 92

- **Oct 2005: Automatic Construction of Multifaceted Browsing Interfaces;** Wisam Dakka, Panagiotis G. Ipeirotis, and Kenneth R. Wood; Proceedings of the 2005 ACM International Conference on Information and Knowledge Management (CIKM), 2005 ; 76/425 = 18% accepted. Citations: 118
- **Apr 2005: Modeling and Managing Content Changes in Text Databases;** Panagiotis G. Ipeirotis, Alexandros Ntoulas, Junghoo Cho, and Luis Gravano; Proceedings of the 21st IEEE International Conference on Data Engineering (ICDE), 2005; 67/521 = 13% accepted; Best Paper Award. Citations: 45
- **Jun 2004: When one Sample is not Enough: Improving Text Database Selection Using Shrinkage;** Panagiotis G. Ipeirotis and Luis Gravano; Proceedings of the 2004 ACM International Conference on Management of Data (SIGMOD), 2004 ; 69/431 = 16% accepted. Citations: 57
- **May 2003: Text Joins in an RDBMS for Web Data Integration;** Luis Gravano, Panagiotis G. Ipeirotis, Nick Koudas, and Divesh Srivastava; Proceedings of the 12th International World-Wide Web Conference (WWW), 2003 ; 13% accepted, more than 600 submissions. Citations: 307
- **Mar 2003: Text Joins for Data Cleansing and Integration in an RDBMS;** Luis Gravano, Panagiotis G. Ipeirotis, Nick Koudas, and Divesh Srivastava; Proceedings of the 19th IEEE International Conference on Data Engineering (ICDE), 2003. Citations: 63
- **Aug 2002: Distributed Search over the Hidden-Web: Hierarchical Database Sampling and Selection;** Panagiotis G. Ipeirotis and Luis Gravano; Proceedings of the 28th International Conference on Very Large Databases (VLDB), 2002 ; 69/432 = 16% accepted. Citations: 221
- **Jul 2002: Extending SDARTS: Extracting Metadata from Web Databases and Interfacing with the Open Archives Initiative;** Panagiotis G. Ipeirotis, Tom Barry, and Luis Gravano; Proceedings of the Second ACM+IEEE Joint Conference on Digital Libraries (JCDL), 2002 ; 33% accepted. Citations: 9
- **Sep 2001: Approximate String Joins in a Database (Almost) for Free;** Luis Gravano, Panagiotis G. Ipeirotis, Hosagrahar Visvesvaraya Jagadish, Nick Koudas, Shanmugaelayut Muthukrishnan, and Divesh Srivastava; Proceedings of the 27th International Conference on Very Large Databases (VLDB), 2001 ; 59/339 = 17% accepted. Citations: 579
- **Jun 2001: SDLIP + STARTS = SDARTS. A Protocol and Toolkit for Metasearching;** Noah Green, Panagiotis G. Ipeirotis, and Luis Gravano; Proceedings of the First ACM+IEEE Joint Conference on Digital Libraries (JCDL), 2001. Citations: 26
- **May 2001: Probe, Count, and Classify: Categorizing Hidden-Web Databases;** Panagiotis G. Ipeirotis, Luis Gravano, and Mehran Sahami; Proceedings of the 2001 ACM International Conference on Management of Data (SIGMOD), 2001 ; 15% accepted. Citations: 201

Workshops

- **Dec 2011: Examining the Impact of Search Engine Ranking and Personalization on Consumer Behavior: Combining Bayesian Modeling with Randomized Field Experiments;** Anindya Ghose, Panagiotis G. Ipeirotis, and Beibei Li; Proceedings of the Workshop on Information Systems and Economics (WISE), 2011. Citations: 2
- **Aug 2011: Whats the Right Price? Pricing Tasks for Finishing on Time;** Siamak Fari-dani, Björn Hartmann, and Panagiotis G. Ipeirotis; Proceedings of the 3rd Human Computation Workshop (HCOMP), 2011; 16/52=30% accepted. Citations: 84
- **Aug 2011: Beat the Machine: Challenging Workers to Find the Unknown Unknowns;** Josh Attenberg, Panagiotis G. Ipeirotis, and Foster Provost; Proceedings of the 3rd Human Computation Workshop (HCOMP), 2011; 16/52=30% accepted

- **May 2011: The Need for Standardization in Crowdsourcing;** Panagiotis G. Ipeirotis and John J. Horton; Proceedings of the Workshop on Crowdsourcing and Human Computation at CHI, 2011. Citations: 21
- **Mar 2011: A Demo Search Engine for Products;** Beibei Li, Anindya Ghose, and Panagiotis G. Ipeirotis; Proceedings of the 20th International World-Wide Web Conference (WWW), 2011 . Citations: 1
- **Feb 2011: Estimating the Completion Time of Crowdsourced Tasks Using Survival Analysis Models;** Jing Wang, Siamak Faridani, and Panagiotis G. Ipeirotis; Proceedings of the Workshop on Crowdsourcing for Search and Data Mining (WSDM11-CSDM), 2011. Citations: 41
- **Jan 2011: Managing Crowdsourcing Workers;** Jing Wang, Panagiotis G. Ipeirotis, and Foster Provost; Winter Conference on Business Intelligence, 2011. Citations: 24
- **Jul 2010: Quality Management on Amazon Mechanical Turk;** Panagiotis G. Ipeirotis, Foster Provost, and Jing Wang; Proceedings of the Second Human Computation Workshop (HCOMP), 2010. Citations: 639
- **Sep 2009: Cramer8217s Rule: How Information Content Moves Markets;** Sinan Aral, Panagiotis G. Ipeirotis, and Sean J. Taylor; Proceedings of the 20th Workshop on Information Systems and Economics (WISE), 2009
- **Sep 2009: The Economic Impact of User-Generated Content on the Internet: Combining Text Mining with Demand Estimation in the Hotel Industry;** Anindya Ghose, Panagiotis G. Ipeirotis, and Beibei Li; Proceedings of the 20th Workshop on Information Systems and Economics (WISE), 2009
- **Oct 2008: Answering General Time Sensitive Queries;** Wisam Dakka, Luis Gravano, and Panagiotis G. Ipeirotis; Proceedings of the 2008 ACM CIKM International Conference on Information and Knowledge Management (CIKM), 2008. Citations: 135
- **Jun 2008: Stay Elsewhere? Improving Local Search for Hotels Using Econometric Modeling and Image Classification;** Beibei Li, Anindya Ghose, and Panagiotis G. Ipeirotis; Proceedings of the Sixth International Workshop on the Web and Databases (WebDB), 2008 ; 14/30 = 46% accepted. Citations: 6
- **Jan 2008: Detecting Important Events Using Prediction Markets, Text Mining, and Volatility Modeling;** George Tziralis and Panagiotis G. Ipeirotis; Proceedings of the Third Workshop on Prediction Markets, 2008
- **Jan 2008: The Impact of Information Disclosure on Stock Market Returns: The Sarbanes-Oxley Act and the Role of Media as an Information Intermediary;** Karthik Balakrishnan, Anindya Ghose, and Panagiotis G. Ipeirotis; Proceedings of the Seventh Workshop on the Economics of Information Security (WEIS), 2008
- **Apr 2007: Faceted Browsing over Large Databases of Text-Annotated Objects;** Wisam Dakka, Panagiotis G. Ipeirotis, and Kenneth R. Wood; Proceedings of the 23rd IEEE International Conference on Data Engineering, Demonstrations (ICDE), 2007 ; 28/73 = 38% accepted. Citations: 6
- **Aug 2006: Automatic Discovery of Useful Facet Terms;** Wisam Dakka, Rishabh Dayal, and Panagiotis G. Ipeirotis; Proceedings of the ACM SIGIR 2006 Workshop on Faceted Search, 2006. Citations: 35
- **Jan 2006: Designing Ranking Systems for Consumer Reviews: The Impact of Review Subjectivity on Product Sales and Review Quality;** Anindya Ghose and Panagiotis G. Ipeirotis; Proceedings of the 2006 Workshop on Information Technology and Systems (WITS), 2006 ; 36/105 = 35% accepted. Citations: 70
- **Aug 2005: Reputation Premiums in Electronic Peer-to-Peer Markets: Analyzing Textual Feedback and Network Structure;** Anindya Ghose, Panagiotis G. Ipeirotis, and Arun

Sundararajan; Proceedings of the ACM SIGCOMM 2005 Third Workshop on Economics of Peer-to-Peer Systems, (P2PEcon), 2005. Citations: 29

- **Jun 2003: Modeling Query-Based Access to Text Databases;** Eugene Agichtein, Panagiotis G. Ipeirotis, and Luis Gravano; Proceedings of the Sixth ACM SIGMOD International Workshop on the Web and Databases (WebDB), 2003 ; 25% accepted. Citations: 32
- **Jun 2001: PERSIVAL Demo: Categorizing Hidden-Web Resources;** Panagiotis G. Ipeirotis, Luis Gravano, and Mehran Sahami; Proceedings of the First ACM+IEEE Joint Conference on Digital Libraries (JCDL), 2001. Citations: 15
- **May 2000: Automatic Classification of Text Databases through Query Probing;** Panagiotis G. Ipeirotis, Luis Gravano, and Mehran Sahami; Proceedings of the Third International Workshop on the Web and Databases, (WebDB), 2000 ; 29% accepted. Citations: 25

Misc publications

- **Jun 2015: The Global Opportunity in Online Outsourcing;** Siou Chew Kuek, Cecilia Paradi-Guilford, Toks Fayomi, Saori Imaizumi, Panos Ipeirotis, Patricia Pina, and Manpreet Singh; The World Bank. textitCitations: 9
- **Mar 2011: The Computer is the New Sewing Machine: Benefits and Perils of Crowdsourcing;** Praveen Paritosh, Panagiotis G. Ipeirotis, Matt Cooper, and Siddharth Suri; Panel at the 20th International World-Wide Web Conference (WWW), 2011
- **Mar 2011: Managing Crowdsourced Human Computation: A Tutorial;** Panagiotis G. Ipeirotis and Praveen K. Paritosh; Tutorial at the 20th International World-Wide Web Conference (WWW), 2011 . textitCitations: 59
- **Dec 2010: System, Method, Software Arrangement and Computer-Accessible Medium for Incorporating Qualitative and Quantitative Information into an Economic Model;** ; US Patent 7,848,979
- **Mar 2010: Improving Product Search with Economic Theory;** Beibei Li, Panagiotis G. Ipeirotis, and Anindya Ghose; Proceedings of the Workshops of 2010 IEEE 26th International Conference on Data Engineering (ICDE), 2010
- **Jan 2009: Searching Digital Libraries;** Panagiotis G. Ipeirotis; Encyclopedia of Database Systems, Part 19, Springer, 2009
- **Jan 2009: Taxonomy Design for Faceted Search;** Wisam Dakka , Panagiotis G. Ipeirotis, and Giovanni Maria Sacco; The Information Retrieval Series, Vol. 25, 175-213, Springer, 2009 . textitCitations: 1
- **Jan 2004: Classifying and Searching Hidden-Web Text Databases;** Panagiotis G. Ipeirotis; Ph.D. Dissertation, Columbia University (advisor: L. Gravano), September 2004. textitCitations: 10

Working Papers

- **Dec 2015: On Hiring Decisions in Online Labor Markets;** Marios Kokkodis, Panagiotis Papadimitriou, and Panagiotis Ipeirotis;
- **Jun 2013: Quality-Based Pricing for Crowdsourced Workers;** Jing Wang, Panagiotis G. Ipeirotis, and Foster Provost; NYU-CBA Working Paper CBA-13-06. Citations: 28
- **Dec 2010: Modeling Dependency in Prediction Markets;** Nikolay Archak and Panagiotis G. Ipeirotis; NYU Center for Digital Economy Research Working Paper No. CEDER-10-05

- **Mar 2010: Demographics of Mechanical Turk;** Panagiotis G. Ipeirotis; NYU Center for Digital Economy Research Working Paper CeDER-10-01. Citations: 372
- **Mar 2009: The Dimensions of Reputation in Electronic Markets;** Anindya Ghose, Panagiotis G. Ipeirotis, and Arun Sundararajan; NYU Center for Digital Economy Research Working Paper No. CeDER-06-02. Citations: 99

SERVICE

- **University Service**
 - University Research Challenge Fund Committee, Provosts Office, New York University
 - IS Recruiting Committee, IOMS Department, New York University
 - Academic Programs Teaching Resources Committee, Stern School of Business, New York University
 - PhD Admissions Committee, IOMS Department, New York University
 - Stern Technology Advisory Committee, New York University
 - Stern PhD Program Task Force, New York University
 - Global Task Force, Stern School of Business, New York University
 - Information Systems / CeDER Research Seminar Coordinator, IOMS Department, New York University
 - NYU Institutional Repository Working Group, New York University
 - PhD Admissions Committee, Student Member, Computer Science Department, Columbia University
 - Board of Directors, Elected Student Representative, Computer Engineering and Informatics Department, University of Patras, Greece
- **Editorial Board, Associate Editor**
 - Information Systems Research, Special Issue on Social Media and Business Transformation
 - INFORMS Journal on Computing (JoC)
 - IEEE Transactions on Knowledge and Data Engineering (TKDE)
 - International Journal of Cooperative Information Systems (IJCIS)
 - ACM Journal on Information and Data Quality (JDIQ)
 - ACM Transactions on Economics and Computation (TEAC), Special Issue with selected papers from ACM 2012
- **Program Committee Chair**
 - Third AAAI Conference on Human Computation Crowdsourcing (HCOMP 2015), co-chaired with Liz Gerber
 - Thirteenth ACM Conference on Electronic Commerce (EC 2012), co-chaired with Kevin Leyton-Brown
 - 3rd Human Computation Workshop (HCOMP 2011), co-chaired with Luis von Ahn
 - Fourth International Workshop on Ranking in Databases (DBRank 2010), co-chaired with Amlie Marian
- **Program Committee Member**
 - 26th International World Wide Web Conference (WWW 2017), Track Chair, Crowdsourcing

- Tenth ACM International Conference on Web Search and Data Mining (WSDM 2017), Senior PC Member
- Workshop on Data Science for Macro-Modeling (DSMM 2016)
- Weaving Relations of Trust in Crowd Work: Transparency and Reputation across Platforms (TRUSTINCW 2016)
- 25th International Joint Conference on Artificial Intelligence (IJCAI-16), Senior PC Member
- 17th ACM Conference on Economics and Computation (EC16) , Senior PC Member
- 25th International World Wide Web Conference, WWW 2016
- Eighth ACM Conference on Web Search and Data Mining (WSDM 2015) , Senior PC Member
- ACM SIGMOD Conference 2015
- ACM Multimedia Conference 2014
- 17th International Workshop on the Web and Databases (WebDB 2014)
- International AAAI Conference on Weblogs and Social Media (ICWSM 2014), Senior PC Member
- 23rd International World Wide Web Conference (WWW 2014), Track Chair, Crowd Phenomena
- Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI-14)
- Fourteenth ACM Conference on Electronic Commerce (EC13), Senior PC Member
- International AAAI Conference on Weblogs and Social Media (ICWSM 2013), Senior PC Member
- ACM SIGMOD Conference 2013, Industrial Track
- 21st International World Wide Web Conference (WWW 2012)
- 31st ACM SIGMOD-SIGACT-SIGART Symposium on Principles Of Database Systems (PODS 2012)
- 20th ACM Conference on Information and Knowledge Management (CIKM 2011) Senior PC Member
- The 17th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2011)
- Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI 2011)
- 20th International World Wide Web Conference (WWW 2011)
- The 34th Annual International ACM SIGIR Conference (SIGIR 2011)
- 14th International Workshop on the Web and Databases (WebDB 2011)
- Crowdsourcing for Search and Data Mining (CSDM 2011)
- 11th ACM Conference on Electronic Commerce (EC 2010)
- The 16th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2010)
- 36th International Conference on Very Large Databases (VLDB 2010)
- The 33rd Annual International ACM SIGIR Conference (SIGIR 2010) Workshops
- The 19th International World Wide Web Conference (WWW 2010)
- 26th IEEE International Conference on Data Engineering (ICDE 2010)
- 13th International Workshop on the Web and Databases (WebDB 2010)
- Social Media Analytics Workshop 2010 (at KDD 2010)

- Human Computation Workshop 2010 (HCOMP 2010) at KDD 2010
- The 10th ACM Conference on Electronic Commerce (ACM EC09) Senior PC Member
- The 32nd Annual ACM SIGIR Conference (ACM SIGIR 2009)
- 35th International Conference on Very Large Data Bases (VLDB 2009)
- The 18th International World Wide Web Conference (WWW 2009)
- Workshop on Web Search Click Data (WSCD 2009)
- Human Computation Workshop (HCOMP 2009)
- 34th International Conference on Very Large Data Bases (VLDB 2008)
- ACM SIGMOD Conference 2008
- The 31st Annual International ACM SIGIR Conference (ACM SIGIR 2008)
- ACM Special Interest Group on Knowledge Discovery and Data Mining (ACM KDD 2008)
- 2008 IEEE 24th International Conference on Data Engineering (ICDE 2008)
- 17th International World Wide Web Conference (WWW 2008)
- Second International Workshop on Ranking in Databases (DBRank 2008)
- 33rd International Conference on Very Large Data Bases (VLDB 2007)
- ACM SIGMOD 2007 (Demonstrations)
- The 30th Annual International ACM SIGIR Conference (ACM SIGIR 2007)
- 9th International Conference on Electronic Commerce (ICEC 2007)
- 1st International Workshop on Ranking in Databases (DBRank 2007)
- Workshop on Text Data Mining and Management (TDMM 2007)
- International Workshop on Dynamic Taxonomies and Faceted Search (FIND 2007)
- 25th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (ACM SIGMOD 2006)
- The 29th Annual International ACM SIGIR Conference (ACM SIGIR 2006)
- 15th International World Wide Web Conference (WWW 2006)
- 22nd International Conference on Data Engineering (IEEE ICDE 2006)
- The 2nd International Workshop on Challenges in Web Information Retrieval and Integration (WIRI 2006)
- International Workshop on Information Quality in Information Systems (IQIS 2006)
- The 28th Annual International ACM SIGIR (ACM SIGIR 2005) Conference on Research and Development in Information Retrieval
- The 14th International World Wide Web Conference (WWW 2005)
- (SPIRE 2005)
- Sixteenth ACM Conference on Hypertext and Hypermedia (ACM Hypertext 2005)
- Workshop on Peer to Peer and Service Oriented Hypermedia (ACM Hypertext 2005 WS4)

- **Ph.D. Dissertation Committees**

- Nikolay Archak, New York University, IOMS Department
- Wisam Dakka, Columbia University, Computer Science Department

- Alpa Jain, Columbia University, Computer Science Department
- Caryn Conley, New York University, IOMS Department
- Gal Oestreicher-Singer, New York University, IOMS Department
- Ziyang Wang, New York University, Computer Science Department

- **Ph.D. Thesis Proposal Committees**

- Beibei Li, New York University, IOMS Department
- Daniel Engler, New York University, Management Department
- Nikolay Archak, New York University, IOMS Department
- Tianyi Jiang, New York University, IOMS Department
- Wisam Dakka, Columbia University, Computer Science Department
- Alpa Jain, Columbia University, Computer Science Department
- Gal Oestreicher-Singer, New York University, IOMS Department
- Caryn Conley, New York University, IOMS Department

- **Organizing Committee**

- ACM SIGMOD International Conference on Management of Data (SIGMOD 2013), Student Sponsorship Chair
- 4th Human Computation Workshop (HCOMP 2012)
- 3rd Human Computation Workshop (HCOMP 2011)
- 2nd Human Computation Workshop (HCOMP 2010)
- 16th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2010)
- 1st Human Computation Workshop (HCOMP 2009)
- Fourth New York Area DB/IR Day (DBIR, Fall 2006)

- **Ad-hoc Journal Referee**

- Management Science
- Marketing Science
- Information Systems Research
- ACM Transactions on the Web
- ACM Transactions on Information Systems
- ACM Transactions on Database Systems
- IEEE Transactions on Knowledge and Data Engineering
- IEEE Transactions on Parallel and Distributed Systems
- Elsevier Information Systems
- INFORMS Journal on Computing
- IEEE Computer
- IEEE Internet Computing

- **Miscellaneous**

- GSTEM Mentor Summer 2014

- External Reviewer, The ACM CHI Conference on Human Factors in Computing Systems (ACM CHI 2011)
- Panelist, NSF Funding Panel
- Panelist, NSF Funding Panel
- Panelist, NSF Funding Panel
- Association for Computing Machinery Special Interest Group on Management of Data (ACM SIGMOD 2005)
- Web Chair, Seventh International Workshop on the Web and Databases (WebDB 2004)

- **External Conference Reviewer**

- 2004 IEEE 20th International Conference on Data Engineering (IEEE ICDE)
- 28th International Conference on Very Large Data Bases (VLDB)
- ACM SIGMOD
- The 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (ACM SIGIR)
- ICDE 2001 -IEEE International Conference on Data Engineering (IEEE ICDE)