

Edward Y. Chang

President, Research & Healthcare, HTC Inc.

E-Mail: eyuchang@gmail.com

URL: <http://infolab.stanford.edu/~echang/>

Education

- 1995-99: Ph.D. Electrical Engineering, Stanford University
- 1992-94: M.S. Computer Science, Stanford University
- 1984-85: M.S. Industrial Engineering and Operations Research, University of California, Berkeley

Professional Employment

- 2015 – Present: President, Corporate Research & Healthcare BU, HTC
- 2012 – 2015: Vice President, Corporate Research & Innovation, HTC
- 2012 – 2015: Adjunct Professor, Computer & Information Science, HKUST, HK
- 2006 – 2012: Director of Research, Google
- 2006 – 2009: Professor, Electrical and Computer Engineering, University of California, Santa Barbara
- 2003 – 2006: Associate Professor (with tenure), Electrical and Computer Engineering, University of California, Santa Barbara
- 2004 – 2006: Chief Technical Advisor, Video Surveillance, Proximex Inc., Cupertino (acquired by ADT)
- 1999 – 2003: Assistant Professor, Electrical and Computer Engineering, University of California, Santa Barbara
- 1998: Teaching Fellow, Stanford University
- 1994 – 1995: Staff Engineer, Corba & JAVA, Sun Microsystems, Mountain View, California
- 1987 – 1994: Principal Software Engineer, Operating System kernels, Digital Equipment Corporation, Palo Alto, California
- 1985 – 1987: Software Engineer, Shop Floor Control Software, Applied Materials, Mountain View, California

Honors and Awards

- XPRIZE Tricorder 2nd place winner with US\$1,000,000 prize, April 2017.
- Fellow of IEEE for contributions to scalable machine learning, 2016.
- Top-10 world most influential scholars in multimedia research, AMiner 2016.
- Best Fitness Tracker Award, HTC/UnderAmor Healthbox, CES 2016.
- Best Technical Demonstration (360 panorama), ACM Multimedia Conference, 2015.
- Best Paper, WWW Conference, 2010.
- Google Innovation Award (for launching Google Q&A to 60+ countries), 2009.
- Best Student Paper, ACM Multimedia Conference, 2003.
- NSF CAREER Award (on small data machine learning), 2002.
- IBM Faculty Partnership Award, 2000, 2001, 2002.

Patents

1. Aristo, an augmented reality platform, with HTC, 2017 (filed).
2. Context-aware symptom checking, with HTC, 2017 (filed).
3. Distributed optimization for speeding up deep learning training, with HTC, 2017 (filed).
4. Data-driven inside-out tracking camera pose estimate, with HTC, 2016 (filed).
5. VIVE paper, a virtual bookstand, with HTC, 2016 (filed).

6. Attention driven optimization techniques for VR/AR systems, with HTC, 2016 (filed).
7. Smart medication reminder, with HTC, 2015 (filed).
8. Bio-signal disease classifier using transfer learning, with HTC (HTC104235-0-US), 2015 (Filed).
9. Classifier for skin disease detection using transfer learning, with HTC (HTC103162-1-US), 2015 (Filed).
10. Searching people via active learning and scene, time, and location fusion, with HTC (HTC103228-1-US), 2014 (Filed).
11. Detecting correct photo-taking gesture for multi-view simulation with single camera, with HTC (HTC102077-0-US), 2014 (Filed).
12. Communicative connection method among multiple devices, with HTC (HTC101679), 2013 (Filed).
13. Method and device for sharing digital object over mesh network, with HTC (HTC101678-0-US), 2013 (Filed).
14. Gyro sensor powered panorama images generation, with HTC (HTC101752-0-US), 2013 (Filed).
15. Mobile electronic device and connection establishment method between mobile electronic devices, with HTC (HTC101554-0-US), 2013 (Filed).
16. Moving direction determination with noisy signals from INS on mobile devices, with Google, 2013 (Pending, US #20,130,158,941).
17. Techniques for feature extraction, with Google, 2017 (US Patent #9,547,914).
18. Dynamic calibration of MEMS sensors, with Google, 2017 (US Patent #9,625,290 B2).
19. Crowd sourcing indoor location, with Google, 2016, US Patent #9,335,175).
20. Barometric pressure sensor based orientation measurement, with Google, 2016 (US Patent, #9,383,202 B2).
21. System and method for attitude correction, with Google, 2016 (US Patent #20,160,033,301).
22. Signal processing to extract pedestrian's moving direction, with Google, 2016 (US Patent, #9,295,027).
23. Context specific language model for input method editor, Google, 2015 (US Patent #90,47,870 B2).
24. Gyroscope calibration, with Google, 2015 (US Patent #20,150,260,545).
25. Hardware attitude detection implementation of mobile devices with MEMS motion sensors, with Google, 2015 (US #9,135.802).
26. P2P location service, with Google, 2013 (US Patent #8,364,172).
27. AdHeat: social network ads model considering influence, with Google, 2013 (US Patent #8,600,812).
28. GPS and MEMS hybrid location-detection architecture, with Google, 2013, (US Patent #8,362,949).
29. Object recognition in images, with Google, 2014 (US Patent #8,798,375).
30. Location monitoring using clusters of entities, with Google, 2014 (US Patent #8,792,905).
31. Ranking user generated content with Google, 2015 (US Patent #8,965,883).
32. Providing advertisements on a social network, with Google, 2012 (US Patent #20,120,226,564).
33. Parallel generation of topics from documents, with Google, 2014 (US Patent #8,898,172).
34. Bidding on users, with Google, 2009 (Pending US #20,100,228,631).
35. Adaptive multi-modal integrated biometric identification detection and surveillance systems, with Proximex, 2004 (US Patent #20,060,093,190).
36. Multimodal high-dimensional data fusion for classification and identification, with Proximex, 2007, (US Patent #7,242,810).
37. Video surveillance using stationary-dynamic camera assemblies for wide-area video surveillance and allow for selective focus-of-attention, with Proximex, 2007 (US Patent #20,060,203,090).
38. Maximizing expected generalization algorithm, with VIMA, 2007 (US Patent #7,158,970).
39. Dynamic partial function in measurement of similarity of objects, with VIMA, 2006 (US Patent #7,106,903).
40. Robust perceptual color identification, with Proximex, 2006 (US Patent #7,136,524).
41. Perception-based image retrieval, with UC, 2005 (US Patent #6,865,302).
42. Computer user interface for perception-based information retrieval, with VIMA, 2003 (US Patent #20,030,016,250).
43. Personalized navigation tree, with NEC, 2002 (US Patent #6,393,427).
44. Client-based interactive digital television architecture, with UC, 2002 (US Patent #20,020,124,259).

45. System for distributed computation processing includes dynamic assignment of predicates to define interdependencies, with DEC, 1994 (US Patent 5,329,626).
46. Rule driven transaction management system and method, with DEC, 1992 (US Patent #69,131,500).

Editorships

- Associate editor, IEEE Transactions on Knowledge and Data Engineering, 2003 – 2007.
- Associate editor, ACM Multimedia Systems Journal, 2005 – 2009.

Program Committees and Reviewers

- Area chair, Deep Learning, ACM Multimedia, 2016.
- Honorable co-chair, IEEE Big Data Congress, Satellite Session, Taipei, 2014.
- Program co-chair, International Joint Conference on Artificial Intelligence (IJCAI), Beijing, 2012.
- Program co-chair, IEEE International Conference on Data Engineering (ICDE), Los Angeles, 2010.
- Program co-chair, WWW, North Carolina, 2010.
- Program co-chair, ACM CIVR, Greece, 2009.
- Program area chair, Deep Learning, ACM Multimedia, Amsterdam, 2016.
- Workshop co-chair, ACM Mobile Location-Based Service (MLDBS), Beijing China, 2011.
- Workshop co-chair, ACM Multimedia, Beijing, 2009.
- Panel program chair, IEEE International Conference on Data Engineering (ICDE), Shanghai China, 2009.
- Program vice-chair, WWW Conference, Beijing China, 2008.
- Program vice-chair, IEEE International Conference on Data Engineering (ICDE), Cancun Mexico, 2008.
- Special session chair, ACM International Conference on Multimedia, Germany, 2007.
- General co-chair, International Conference on Multimedia Modeling, Beijing, January 2006.
- General co-chair, ACM International Workshop on Video Surveillance and Sensor Networks, November 2005.
- Conference co-chair, IS&T/SPIE Conf. on Storage and Retrieval Methods and Applications for Multimedia, 2005, 2006.
- Program co-chair, ACM Workshop on Video Surveillance and Sensor Networks, 2003, 2004.
- Program co-chair, IEEE International Workshop on Managing Data for Emerging Multimedia Applications, 2005.
- Co-chair, ARDA Workshop on A Large Scale Concept Ontology for Multimedia Understanding, 2004.
- Panel co-chair, ACM International Conference on Multimedia, Singapore, November 2005.
- Corporate Sponsor chair, ACM SIGMOD Conference, 2007.
- Area chair, IEEE Conference on Multimedia, 2004, 2005.
- Panelist, National Science Foundation, 2002 – 2009.
- Senior program committee member of WSDM, 2012.
- Program committee member of ACM KDD, 2007, 2010, 2011, 2012, 2014, 2015, 2016.
- Program committee member of AAAI, 2004, 2014, 2015, 2016, 2018.
- Program committee member of VLDB, 2006, 2008, 2009, 2010.
- Program committee member of WWW 2006, 2008, 2009, 2010.
- Program committee member of ACM Multimedia Conf., 2002-2011, 2013, 2016.
- Program committee member of the ACM SIGMOD Conf., 2004.
- Program committee member of the ACM CIKM Conf., 2003, 2004, 2008, 2009, 2010.
- Program committee member of the SIAM Data Mining Conf., 2004.
- Program committee member of the IEEE Big Data Conf., 2015, 2016.
- Program committee member of the IEEE Multimedia Conf., 2002, 2003, 2004.
- Program committee member of the IEEE Data Engineering Conf. (ICDE), 2002, 2007, 2009.
- Program committee member of International Conference on Computer Vision (ICCV), 2003.

- Local arrangement chair of the ACM SIGMOD Conf., 2001.

Selected Keynotes, Tutorials, Panels, and Invited Talks

1. DeepQ: Advancing Healthcare through AI and VR (**keynote**), IEEE International Conference on Multimedia Information Retrieval and Processing (MIPR), Miami, April 2018.
2. DeepQ: Advancing Healthcare through AI and VR (**keynote**), XXI Symposium Neuroradiologicum, Taipei, March 2018.
3. AI for Cure, TEDx Taiwan, December 2017.
4. DeepQ: Advancing Healthcare through AI and VR (**keynote**), The ACM International Conf. on Multimedia, Mountain View, October 2017.
5. Developing AlphaGo for Powering Healthcare & VR (**distinguished lecture**), Tsinghua University, Beijing, March 2017.
6. Developing AlphaGo for Powering Healthcare & VR (**keynote**), The ACM/IEEE International Conf. on Microarchitectures, October 2016.
7. Developing AlphaGo for Powering Healthcare & VR (**keynote**), The IEEE International Conf. on Big Multimedia Data Conference, April 2016.
8. Not-So-Big Clinical Data Predictive Analytics, (**keynote**), The 24th Wireless Optical and Communication Conference, Taipei, October 2015.
9. Predictive Analytics in Healthcare (**panelist**), Stanford Data Science Initiative Annual Retreat, Stanford University, October 2015.
10. Big Data Analytics, Architectures, Algorithms, and Applications (**keynote**), The IEEE Big Data Congress, Taipei, May 2015.
11. Signal Fusion and Big Data Analytics on Massive Sensor Data Sets (**distinguished lecture**), HKUST, April 2015.
12. Signal Fusion and Big Data Analytics on Massive Sensor Data Sets (**keynote**), The IEEE 10th International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP), Singapore, April 2015.
13. Big Data Analytics, Architectures, Algorithms, and Applications (**invited tutorial**), BigDat 2015, Tarragona Spain, January 2015.
14. Data Mining in Big Data Era: Opportunities, Challenges and Practice (**Panelist**), PAKDD, May 2014.
15. Big Data and Bio Informatics (**keynote**), The 4th International Symposium on Dynamical Biomarkers for Translational Medicine, Taipei, April 2014.
16. Making Smartphones Smarter (**Industry keynote**), [ACM MobiSys](#), Taipei, June 2013.
17. Mobile Information Management and Retrieval, The Next Frontier (**keynote**), [ASIST 75th Annual Meeting](#), Baltimore, October 2012.
18. Big Data Analysis, Research Issues (**invited talk**), China National Academy of Science, May 2012.
19. From Machine Learning to Human Innovation (**invited lecture**), School of Engineering, Lanzhou University, Lanzhou, China, May 2012.
20. From Machine Learning to Human Innovation (**keynote**), City U. of HK Media Center Grand Opening, Hong Kong, March 2012.
21. Mobile Information Management and Retrieval (**invited talk**), CS Seminar, Taiwan University, January 2012.
22. XINS --- An Anatomy of Indoor Positioning and Navigation Architecture, Google Faculty Summit (**keynote**), November 2011.
23. Mobile Information Management, International Conference on Asia-Pacific Digital Libraries (**keynote**), Tsinghua U., October 2011.
24. Sensor-Aided Mobile Information Management and Retrieval, ACM SIGIR (**keynote**), Beijing, 2011.
25. Search and Social Integration, The 11th International Conf. on Web Information System Engineering (**keynote**), Hong Kong, December 2010.
26. Developing Next Generation Cyber-Infrastructure for Mobile IT Revolution (**keynote**), Google Faculty Summit, Shanghai, November 2010.

27. Processing Web-Scale Multimedia Data (**tutorial**), ACM International Conference on Multimedia, Florence, October 2010.
28. Scalable Algorithms and Systems for Mining Massive Datasets, Computer and Information Science, ETH, Zurich, October 2010.
29. Confucius and "its" Intelligent Disciples (**keynote**), the 27th National Database Conference, Beijing, October 2010.
30. AdHeat, A New Influence-based Social Ads Model and its Tera-Scale Algorithms (**invited tutorial**) [MMDS](#), Stanford University, June 2010.
31. AdHeat, A New Influence-based Social Ads Model and its Tera-Scale Algorithms (**keynote**) UWAP, Hawaii, June 2010.
32. Confucius and "its" Intelligent Disciples (**keynote**), the 18th ACM International Conference on Information and Knowledge Management (CIKM), Hong Kong, November 2009.
33. Parallel Algorithms for Mining Large-scale Multimedia Datasets (**tutorial**), ACM International Conference on Multimedia, October 2009.
34. Confucius and "its" Intelligent Disciples (**keynote**), the 5th International Conference on Advanced Data Mining and Applications, Beijing, August 2009.
35. From Machine Learning to Human Innovation (**invited talk**), IEEE Student Leadership Meeting, Singapore, July 2009.
36. Parallel Computing Paradigms for Machine Learning Algorithms (**invited talk**), National Science Foundation Workshop, Copenhagen, July 2009.
37. Large-scale Collaborative Filtering (**keynote**), The 5th AAIM Conference, San Francisco, June 2009.
38. Large-scale Photo Annotation Using Collective Wisdom of Data and Users (**keynote**), IEEE International Conference on Multimedia Computing and Systems (ICMCS), Ouarzazate, Morocco, April 2009.
39. Parallel Algorithms for Mining Large-Scale Data (**keynote**), Workshop of Modeling, Mining and Managing Evolving Social Networks (in conjunction with ICDE), Shanghai, March 2009.
40. Large-scale Collaborative Filtering (**Andrew Yao lecture**), Tsinghua University, March 2009.
41. From Machine Learning to Human Innovation (**distinguished lecture**), Public Lecture on Mathematics, Hong Kong, Feb. 2009.
42. Large-scale Machine Learning Algorithm (**invited lecture**), Symposium of Machine Learning and Bioinformatics, Hong Kong, Feb. 2009.
43. Large-scale Machine Learning Algorithms (**invited talk**), National Taiwan University/Academia Sinica, Taipei, December 2008.
44. Beyond Search---Computational Intelligence for the Web (**invited talk**), NIPS, Vancouver/Whistler, December 2008.
45. From Machine Learning to Human Innovation (**invited talk**), Annual Meeting of Machine Learning and Applications, Nanjing, November 2008.
46. Large-scale Collaborative Filtering (**invited talk**), UC Berkeley EECS/Math Seminar, Berkeley, October 2008.
47. Social Network Open Platform Strategies and Applications (**invited talk**), CSDN, Shanghai, September 2008.
48. Organizing Multimedia Data Socially with Scalable Algorithms (**keynote**), ACM International Conference on Image and Video Retrieval, Niagara Falls, July 2008.
49. Large-scale Social-Graph Mining, Challenges and Scalable Solutions (**tutorial**), [MMDS](#), Workshop on Algorithms for Modern Massive Data Sets, Stanford University, June 2008.
50. Rich Media and Web 2.0 (**panel moderator**), WWW, Beijing, April 2008.
51. Massive Mining on Social Graphs (**keynote**), Social Network Data Mining Workshop, WWW, Beijing, April 2008.
52. Organizing World's Information, Socially (**invited talk**), Stanford InfoSeminar, Computer Science Department, Stanford University, March 14th, 2008.
53. Parallel, Combinational Collaborative Filtering (**invited talk**), HP Labs, Palo Alto, March 2008.
54. Organizing World's Information, Socially (**invited talk**), Chinese University Hong Kong, Hong Kong, December 2007.

55. Media Sharing on Social Networks (**keynote**), Pacific-rim Multimedia Conference, Hong Kong, December 2007.
56. Web-scale Multimedia Data Management: Challenges and Remedies (**keynote**), Visual and Multimedia Digital Libraries, sponsored by European Commission Networks of Excellence, Modena, Italy, September 2007.
57. Challenges of and Remedies for Large-scale Multimedia Information Retrieval (**keynote**), CVDB in conjunction with ACM SIGMOD, Beijing, June 2007.
58. Internet, Past and Future (**panel moderator**), Internet+ Conference, Beijing, March 2007.
59. Web 2.0 and Multimedia (**keynote**), International Conf. on Multimedia Modeling, Singapore, January 2007.
60. Parallel Support Vector Machines, Google Tech Talk, Mountain View, December 2006.
61. Unified & Scalable Learning for Multimedia Information Retrieval (**keynote**), MIR Workshop at ACM MM Conf., Santa Barbara, October 2006.
62. Web 2.0, Synergy and Challenge (**panelist**), ACM MM Conf., Santa Barbara, October 2006.
63. Statistical Foundations for Bridging the Semantic Gap (**tutorial**), ACM International Conference on Multimedia, Singapore, November 2005.
64. Event Sensing on Distributed Video-Sensor Networks (**keynote**), Basenets 2004, in cooperation with ACM/IEEE Conf. on Broadband Networks, San Jose, October 2004.
65. Machine Learning for Information Retrieval, John Hopkins Applied Physics Lab., April, 2004.
66. Perception-based Image Retrieval, An Interdisciplinary System of Computer Vision, Machine Learning and Databases, UIUC Database Seminar, April 2004.
67. Perception-based Image Retrieval, Google, Mountain View, January 2004.
68. Statistical Learning under Extreme Constraints, National University of Singapore, Singapore, December 2003.
69. High-performance Pre-emptible and MEMS-based IOs, IBM Almaden Research Center, San Jose, November 2003.
70. Statistical Learning under Extreme Constraints, Mitsubishi Research Lab., Cambridge, August 2003.
71. Statistical Learning for Visual Information Analysis and Retrieval (**tutorial**), IEEE International Conference on Multimedia, Baltimore, July 2003.
72. Statistical Learning Methods for Emerging Database Applications (**plenary tutorial**), DASFAA, Kyoto University, Kyoto, March 2003.
73. Multimedia Indexing: Promises and Problems, Panel, IEEE International Conference on Multimedia, Lausanne, Switzerland, August 2002.
74. Research Issues in Multimedia Databases, Research Seminar, UCLA, August 2002.
75. Learning and Measuring Perceptual Similarity, ICASSP Special Session in Statistical Learning Methods, Orlando, May 2002.
76. Interactive TV Infrastructures, SONY Research, San Jose, April 2002.
77. Dynamic Partial Function --- A Perceptual Distance Function for Measuring Similarity, Digital Library Seminar, Computer Science, UC Berkeley, Feb. 2002.
78. Multimedia Data Mining, HP Data Mining Lab., Palo Alto, January 2002.
79. Characterizing, Indexing, and Retrieving Art Imagery, NSF Workshop, Harvard University, November 2001.
80. Measuring and Learning Perceptual Similarity, NSF Workshop, INRIA-Rocquencourt, France, September 2001.
81. Learning Query Concepts via Intelligent Sampling, IBM Almaden Research Center, September 2001.
82. Personalizable Interactive Digital VCR, SONY Research, April 2001.
83. On Managing Continuous Media Data, Compaq Western Research Lab, Palo Alto, September 2000.
84. Personalizable Interactive Digital VCR, IBM T.J. Watson, February 2000.

Publications

Invited Papers

1. Web-scale Image Annotation, J. Liu, R. Hu, et al., PCM Special Session, Tainan, December 2008.
2. On Usage Models of Content-based Image Search, Filtering, and Annotation, D. T.L., E. Y. Chang, K.-T. Cheng, and C.-W. B. Chang, IS&T/SPIE Conference on Internet Imaging, San Jose, January 2006.
3. A Unified Learning Paradigm for Large-scale Personalized Information Management, IEEE Emerging Information Technology Conference, Taipei, August 2005.
4. Multimodal Kernel Fusion for News Video Concept Detection, Y. Wu, C.-Y. Lin, E. Y. Chang and J. R. Smith, IEEE International Conference on Image Processing (ICIP), Singapore, 2004.
5. Active Learning and its Scalability for Image Retrieval, E. Y. Chang and W.-C. Lai, IEEE International Conference on Multimedia (ICME), Taipei, June 2004.
6. Toward Building a Robust and Intelligent Video Surveillance System: A Case Study, E. Y. Chang, Y.-F. Wang, and I.-J. Wang, IEEE International Conference on Multimedia (ICME), Taipei, June 2004.
7. SFINX: A Multi-source Fusion and Mining System, Z. Dimitrijevic, G. Wu, and E. Chang, IEEE Pacific-Rim Conference on Multimedia (PCM), Singapore, December 2003.
8. Multimedia Web Services for Filtering, Searching, and Digital Rights Management, Y.-L. Wu, C.-W. Chang, W.-C. Lai, K.-T. Cheng, and E. Chang, IEEE Pacific-rim Conference on Multimedia (PCM), Singapore, December 2003.
9. Statistical Learning for Effective Visual Information Retrieval, E. Chang, B. Li, G. Wu, and K.-S. Goh, IEEE International Conference on Image Processing (ICIP), Barcelona, September 2003.
10. On Learning Perceptual Distance Functions for Image Retrieval, E. Chang and B. Li, IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Orlando, May 2002.
11. Indexing Multimedia Data in High dimensional and Weighted Feature Spaces, K. Goh and E. Chang, The 6th Visual Database Conference (VDB), Australia, May 2002.
12. Learning and Measuring Perceptual Similarity, E. Chang, NSF/INRIA/Berkeley/IBM MMCBIR Workshop, INRIA - Rocquencourt, France, September 2001.

Books

- Big Data Analytics for Large-Scale Multimedia Search, Wiley & Sons, with Stefanos Vrochidis, Yiannis Kompatsiaris, and Benoit Huet, (expect to be published in Q4 2017).
- Foundation of Large-scale Multimedia Information Management and Retrieval, Springer, E. Y. Chang, 2011.
- Nomadic Eternity (Poetry), E. Y. Chang, 2011.

Book Chapters

1. Representation Learning on Large and Small Data, Chun-Nan Chou, Chuenkai Shie, Mark Chang, J. Chang, and Edward Y. Chang, A book chapter in Big Data Analytics for Large-Scale Multimedia Search, 2017.
2. PSVM: Parallel Support Vector Machines with Incomplete Cholesky Factorization, Edward Y. Chang, Hongjie Bai, Kaihua Zhu, Hao Wang, Jian Li, and Zhihuan Qiu, A book chapter in Scaling up Machine Learning, Cambridge Press, 2011.
3. Large-scale Spectral Clustering with MapReduce and MPI, Wen-Yen Chen, Yangqiu Song, Hongjie Bai, Chih-Jen Lin, and E. Y. Chang, A book chapter in Scaling up Machine Learning, Cambridge Press, 2011.
4. Multimedia Data Semantic Modeling and Knowledge Representation, E. Y. Chang, An invited chapter in Encyclopedia of Database Systems, 2009.

5. Statistical Models and Algorithms for Image Annotation and Retrieval: The Mathematics of Perception E. Y. Chang, A book chapter in Digital Multimedia Perception, 2006.
6. Large-Scale Image Search on the Web, W.-C. Lai, E. Y. Chang and K.-T. Cheng, A book chapter in Web Document Analysis: Challenges and Opportunities, World Scientific – Series Machine Perception and Artificial Intelligence, 2002.
7. Managing Parallel Disks for Continuous Media Data, E. Y. Chang, C. Li and H. Garcia-Molina, A Book Chapter in Information Organization & Databases, p.107-120, Kluwer Publisher, 2000.

Refereed Journal Publications

1. Joint Representation Learning for Location-Based Social Networks with Multi-Grained Sequential Contexts, W. Zhao, F. Fan, J.-R. Wen, and Edward Y. Chang, ACM Transactions on Knowledge Discovery from Data (TKDD), 2017 (accepted).
2. A Probabilistic Lifestyle-Based Trajectory Model for Social Strength Inference from Human Trajectory Data, W. Zhao, N. Zhou, W. Zhang, J.-R. Wen, and Edward Y. Chang, ACM Transactions on Information Systems (TOIS), 35(1) 8, 2016.
3. Connecting Social Media to E-Commerce: Cold-Start Product Recommendation using Microblogging Information, Wayne Xin Zhao, Sui Li, Yulan He, Edward Y. Chang, Ji-Rong Wen, and Xiaoming Li, IEEE Transactions on Knowledge and Data Engineering (TKDE) 28(5): 1147-59, 2016.
4. Mining Product Adopter Information from Online Reviews for Improving Product Recommendation, X. Zhao, J. Wang, Y. He, J-R Wen, E. Y. Chang, and X. Li, ACM Transactions on Knowledge Discovery from Data (TKDD) 10(3), 2016.
5. A Data-Drive Study on Image Feature Extraction and Fusion, Z. Wang, F. Li, E. Y. Chang, and S. Yang, Information Sciences, Elsevier, March 2014.
6. Web-Scale Multimedia Processing and Applications [Scanning the Issue], Edward Y. Chang, Shih-Fu Chang, Alexander G. Hauptmann, Thomas S. Huang, Proceedings of the IEEE 100(9): 2580-2583, 2012
7. An Autoregressive Approach to Nonparametric Hierarchical Dependent Modeling, Zhihua Zhang, Dakan Wang, Edward Y. Chang, Journal of Machine Learning Research Proceedings Track 22: 1416-1424, 2012
8. PLDA+: Parallel Latent Dirichlet Allocation with Data Placement and Pipeline Processing, Zhiyuan Liu, Yuzhou Zhang, Edward Y. Chang, and Maosong Sun, ACM Transactions on Intelligent Systems and Technology 2(3), 2011.
9. Parallel Spectral Clustering in Distributed Systems, Wen-Yen Chen, Yangqiu Song, Hongjie Bai, Chih-Jen Lin, and E. Y. Chang, IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 33(3), 2011.
10. Active Learning for Interactive Multimedia Retrieval, Huang, T. S., Dagli, C. K., Rajaram, S., Chang, E. Y., Mandel, M. I., Poliner, G. E., Ellis, D. P. W., Proceedings of the IEEE, 96(4), pp. 648-67, April 2008.
11. Sliced Coordinate Analysis for Effective Dimension Reduction and Nonlinear Extensions, Zhihua Zhang, J. Kwok, D.-Y. Yeung, and Edward Y. Chang, Journal of Computational & Graphical Statistics (JCGS), Volume 17, Number 1, March 2008.
12. Scalable landmark recognition using EXTENT, Arun Qamra and E. Y. Chang, Multimedia Tools Appl. 38(2): 187-208, 2008.
13. Data Management Projects at Google, Michael Cafarella, E. Y. Chang, Andrew Fikes, Alon Halevy, Wilson Hsieh, Alberto Lerner, Jayant Madhavan, Yossi Matias, and Muthu Muthukrishnan, SIGMOD Record, 2008.
14. Nonparametric Regression Using Reproducing Kernel Hilbert Spaces, Zhihua Zhang, Gang Wu and Edward Y. Chang, IEEE Transactions on Neural Networks, September 2007.
15. Building MEMS-based Storage Systems for Streaming Media, Raju Rangaswami, Zoran Dimitrijevi., Edward Chang, Klaus Schauer, ACM Transactions on Storage 3(2), June, 2007.
16. KDX: An Indexer for Support Vector Machines, N. Panda and E. Y. Chang, IEEE Transactions on Knowledge and Data Engineering (TKDE), 2006.

17. Active Learning in Very Large Image Databases, Navneet Panda, Kingshy Goh, and E. Y. Chang, *Journal of Multimedia Tools and Applications Special Issue on Computer Vision Meets Databases*, 2006.
18. One, Two Class SVMs for Multi-class Image Annotation, K.-S. Goh and E. Y. Chang, *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 17(10), October 2005.
19. Semantics and Feature Discovery via Confidence-based Dynamic Ensemble, K.-S. Goh, B. Li, and E. Y. Chang, *ACM Transactions on Multimedia*, 1(2), pp. 168-189, June, 2005.
20. Support Vector Machine Concept-Dependent Active Learning for Image Retrieval, E. Y. Chang, S. Tong, K.-S. Goh, and C.-W. Chang, *IEEE Transactions on Multimedia*, 2005.
21. Systems Support for Preemptive Disk Scheduling, Z. Dimitrijevic, R. Rangaswami, and E. Y. Chang, *IEEE Transactions on Computers*, October 2005.
22. KBA: Kernel Boundary Alignment Considering Imbalanced Data Distribution, G. Wu and E. Y. Chang, *IEEE Trans on Knowledge and Data Engineering (TKDE)*, 17(6), pp.786-695, May 2005.
23. Enhanced Perceptual Distance Functions and Indexing for Image Replica Recognition, A. Qamra, Y. Meng, and E. Y. Chang, *IEEE Trans on Pattern Analysis and Machine Intelligence (PAMI)*, 27(3), pp.379-391, March 2005.
24. The Anatomy of a Multi-camera Security Surveillance System Special Issue on Video Surveillance, L. Jiao, Y. Wu, G. Wu, E. Y. Chang, Y.-F. Wang, *ACM Multimedia System Journal*, 2004.
25. Disk-aware Data Management for Interactive Media Services, R. Rangaswami, Z. Dimitrijevic, E. Chang, and S.-H. G. Chan, *IEEE Transactions on Multimedia*, Volume 5, Number 4, pp.558-569, December 2003.
26. MEGA --- The Maximizing Expected Generalization Algorithm for Learning Complex Query Concepts, E. Chang and B. Li, *ACM Transactions on Information Systems (TOIS)*, Volume 21, Issue 4, pp.347-382, October 2003.
27. The Discovery of A Perceptual Distance Function for Measuring Image Similarity, B. Li, E. Chang, *ACM Multimedia Journal Special Issue on Content-Based Image Retrieval*, Volume 8, Number 6, pp. 512-522, April 2003.
28. CBSA: Content-based Soft Annotation for Multimodal Image Retrieval Using Bayes Point Machines, E. Chang, K. Goh, G. Sychay, and G. Wu, *IEEE Transactions on Circuits and Systems for Video Technology Special Issue on Conceptual and Dynamical Aspects of Multimedia Content Description*, Volume 13, Number 1, pp.26-38, January 2003.
29. Clindex: Approximate Similarity Queries in High-Dimensional Spaces, C. Li, E. Chang, Hector Garcia-Molina, and Gio Wiederhold, *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, Volume 14, Number 4, pp.792-808, July/August 2002.
30. Answering Queries with Useful Bindings, C. Li and E. Chang, *ACM Transactions on Databases (TODS)*, Volume 26, Number 3, pp. 313-343, September 2001.
31. Maximizing QoS for Interactive DTV Clients, Edward Chang, *The Computer Communications Journal (Special Issue)*, Elsevier, Volume 23, Number 3, February 2000.

Refereed Conference Publications

1. Distributed Training Large-Scale Deep Architectures, Shang-Xuan Zou, Chun-Yen Chen, Chun-Yen Chen, Jui-Lin Wu, Chun-Nan Chou, Chia-Chin Tsao, Kuan-Chieh Tung, Ting-Wei Lin, Cheng-Lung Sung, and Edward Y. Chang (submitted), June 2017.
2. Artificial Intelligence in XPRIZE DeepQ Tricorder, Edward Y. Chang, M-H Wu, K-F Tang, H-C Kao, and C-N Chou (keynote), *ACM International Conference on Multimedia*, Mountain View, October 2017.
3. Aristo, An Augmented Reality Platform for Immersion and Interactivity, ZhongYang Zheng, et al., *ACM International Conference on Multimedia*, Mountain View, October 2017.
4. CLKN: Cascaded Lukas-Kanade Networked for Image Alignment, CheHang Chang, Chun-Nan Chou, and Edward Y. Chang, *IEEE CVPR*, July 2017.
5. Tweet Timeline Generation with Determinantal Point Processes, Jin-ge Yao, Feifan Fan, Wayne Xin Zhao, Xiaojun Wan, Edward Chang, and Jianguo Xiao, *Proceedings of AAAI*, January 2016.

6. Transfer Representation Learning for Medical Image Analysis, Chuenkai Shie, Chung-Hisang Chuang, Chun-Nan Chou, Meng-Hsi Wu, and Edward Chang, IEEE Engineering in Medicine and Biology Society, Milan, August 2015.
7. Network Representation Learning with Rich Text Information, Cheng Yang, Zhiyuan Liu, Deli Zhao, Maosong Sun, and Edward Y. Chang, International Joint Conference on Artificial Intelligence (IJCAI), Buenos Aires, July 2015.
8. Gaussian Processes for High-Dimensional Regression: A Method Based on Deep Neural Networks, Wenbing Huang, Deli Zhao, Fuchun Sun, Huaping Liu, and Edward Y. Chang, International Joint Conference on Artificial Intelligence (IJCAI), Buenos Aires, July 2015.
9. A New Retraction for Accelerating the Riemannian Three-Factor Low-Rank Matrix Completion Algorithm, Zhizhong Li, Deli Zhao, Zhouchen Lin, and Edward Y. Chang, IEEE CVPR, Boston, June 2015.
10. Exact Recoverability of Robust PCA via Outlier Pursuit with Tight Recovery Bounds Hongyang Zhang, Zhouchen Lin, Chao Zhang, and Edward Y. Chang, Proceedings of AAAI, January 2015.
11. Big Data, Small Footprint: The Design of a Low-Power Classifier for Detecting Transportation Modes, M. Yu, T. Yu, C.-J. Lin, and E. Y. Chang, VLDB, Hangzhou, August 2014.
12. Distant Supervision for Relation Extraction with Matrix Completion, M. Fan, D. Zhao, Q. Zhou, Z. Liu, T. F. Zheng, and E. Y. Chang, ACM ACL, June 2014.
13. Context-Aware Computing: Opportunities and Open Issues, E. Y. Chang, VLDB, 2013.
14. MOIST: A Scalable and Parallel Moving Object Indexer with School Tracking, VLDB, Istanbul, August 2012.
15. An Autoregressive Approach to Nonparametric Hierarchical Dependent Modeling, Zhihua Zhang, Dakan Wang, Edward Y. Chang, Journal of Machine Learning Research Proceedings Track 22: 1416-1424, 2012.
16. A Data-Driven Approach to Question Subjectivity Identification in Community Question Answering, Tom Chao Zhou, Xiance Si, Edward Y. Chang, Irwin King, and Michael R. Lyu, AAAI 2012.
17. Error Modeling and Estimation Fusion for Indoor Localization, Weipeng Zhuo, Bo Zhang, Shueng-Han Gary Chan, Edward Y. Chang, p.741-746, ICME 2012.
18. Sublinear Algorithms for Penalized Logistic Regression in Massive Datasets, Haoruo Peng, Zhengyu Wang, Edward Y. Chang, Shuchang Zhou, and Zhihua Zhang, 553-568, ECML/PKDD (1) 2012.
19. Entity Disambiguation with Freebase. Web Intelligence, Zhicheng Zheng, Xiance Si, Fangtao Li, Edward Y. Chang, and Xiaoyan Zhu, 82-89, 2012.
20. Question Identification on Twitter, B. Li, X. Si, E. Y. Chang, I. King, and M. Lyu, ACM CIKM, 2011.
21. Summarization of Yes/No Questions Using a Feature Function Model, J. He, D. Dai, and E. Y. Chang, Asian Conf. on Machine Learning, Tainan, November 2011.
22. K2Q: Generating Natural Language Questions from Keywords with User Refinements, Z. Zheng, X. Si, E. Y. Chang, and X. Zhu, The 5th International Joint Conference on Natural Language Processing (INCNLP), Thailand, November 2011.
23. AdHeat: An Influence-based Diffusion Model for Propagating Hints to Match Ads (**best paper**), H.J. Bao and E. Y. Chang, WWW 2010, April 2010.
24. Confucius and Its Intelligent Disciples, Search + Social, X. Si, E. Y. Chang, and Z. Gyongyi, VLDB, pp. 1505-16, Singapore, Aug. 2010.
25. PLDA---Parallel Dirichlet Allocation for Large-scale Applications, AAIM, San Francisco, June 2009.
26. Discovery of Community User Latent Behavior, Wen-Yen Chen, et al., WWW 2009, Spain, April 2009.
27. PFP: Parallel FP-Growth for Query Recommendation, H. Li, Y. Wang, D. Zhang, M. Zhang, and E. Y. Chang, ACM Recomm. Systems, Lausanne, October 2008.
28. CCF: Combinational Collaborative Filtering for Personalized Community Recommendation, Wen-Yen Chen, Dong Zhang, and E. Y. Chang, ACM KDD, 2008.
29. Parallel Spectral Clustering, Yangqiu Song., Wen-Yen Chen., Chih-Jen Lin, Edward Y. Chang. and Hongjie Bai, ECML, Belgium, September 2008.
30. Organizing Multimedia Data Socially, E. Y. Chang, ACM CIVR, Canada, July 2008.
31. Parallelizing Support Vector Machines on Distributed Computers, Edward Chang, Kaihua Zhu, Hao Wang, Hongjie Bai, Jian Li, Z. Qiu, and Hang Cui, NIPS 2007.

32. Optimizing Matrix Decomposition (invited), Kaihua Zhu, H. Cui, H. Bai, J. Li, Z. Qiu, H. Wang, H. Xu, and Edward Y. Chang, IEEE International Conf. on Multimedia, Beijing, July 2007.
33. Bayesian Reasoning for Sensor Group-Queries and Diagnosis, Ankur Jain, E. Y. Chang, and Y.-F. Wang, DASFAA, April 2007.
34. RCMAP: Efficiently Creating High-Quality Euclidean Embeddings, A. Qamra and E. Y. Chang, SIAM International Conference on Data Mining (SDM), Minneapolis, Minnesota, April 2007.
35. Adaptive Nonlinear Clustering in Data Streams, Ankur Jain, Zihua Zhang, and Edward Y. Chang, ACM International Conference on Information and Knowledge Management (CIKM), November 2006.
36. Mining Blog Stories Using Community-based and Temporal Clustering, Arun Qamra, Belle Tseng, E. Y. Chang, ACM International Conference on Information and Knowledge Management (CIKM), November 2006.
37. Efficient Top-k Hyperplane Query Processing for Multimedia Information Retrieval, N. Panda and E. Y. Chang, ACM International Conference on Multimedia (MM), October 2006.
38. A Distributed Photo Sharing and Searching System, B. Lee, W.-Y. Chen, and E. Y. Chang, ACM International Conference on Multimedia (MM), October 2006.
39. Hypersphere Indexer, Navneet Panda and E. Y. Chang, 17th Conference on Database and Expert Systems Applications (DEXA), Poland, September 2006.
40. Learning the Unified Kernel Machines for Classification, Steven Choi, E. Y. Chang, and M. Lyu, ACM International Conference on Knowledge Discovery and Data Mining (KDD), August 2006.
41. Incremental Approximate Matrix Factorization for Speeding up Support Vector Machines, G. Wu, E. Y. Chang, Y.-K. Chen, and C. Hughes, ACM International Conference on Knowledge Discovery and Data Mining (KDD), August 2006.
42. Concept Boundary Detection for Speeding up SVMs, N. Panda, E. Y. Chang, and G. Wu, International Conference on Machine Learning (ICML), 2006.
43. Color In Motion, G. Wu, A. Rahimi, K.-S. Goh, C. Tsai, A. Jain, Y. Wu, E. Y. Chang, and Y.-F. Wang, IEEE International Conf. on Computer Vision and Pattern Recognition (CVPR), June 2006.
44. Formulating Context-dependent Similarity, Gang Wu, Navneet Panda, Edward Y. Chang, ACM International Conference on Multimedia (MM), Singapore, November 2005 (16% accepted).
45. Multimodal Metadata Fusion Using Causal Strength, Yi Wu, Edward Y. Chang, Belle Tseng, ACM International Conference on Multimedia (MM), Singapore, November 2005 (16% accepted).
46. Learning with Non-metric Proximity Matrices, G. Wu, E. Y. Chang, and Z. Zhang, ACM International Conf. on Multimedia (MM), Singapore, November 2005.
47. Formulating Distance Functions via the Kernel Trick, G. Wu, E. Y. Chang, and N. Panda, ACM International Conference on Knowledge Discovery and Data Mining (KDD), Chicago, August 2005.
48. Manifold Learning: A Promised Land or Work in Progress? M.-C. Yeh, I.-H. Lee, G. Wu, Y. Wu and E. Y. Chang, IEEE International Conference on Multimedia (ICME), Amsterdam, July 2005.
49. Inferring Image Metadata From Context and Content, C. Tsai, A. Qamra, and E. Y. Chang, IEEE International Conf. on Multimedia (ICME), Amsterdam, 2005.
50. Exploiting Geometric Property for Support Vector Machine Indexing, N. Panda, and E. Y. Chang, SIAM International Conference on Data Mining (SDM), Newport Beach, April 2005.
51. Kronecker Factorization for Speeding up Kernel Machines, G. Wu, Z. Zhang, and E. Y. Chang, SIAM International Conf. on Data Mining (SDM), April 2005.
52. Concept-dependent Multimodal Active Learning for Image Retrieval, K. Goh, E. Y. Chang, and W.-C. Lai, ACM International Conference on Multimedia (MM), pp.564-571, New York, October 2004.
53. Optimal Multimodal Fusion for Multimedia Data Analysis, Y. Wu, E. Y. Chang, K. C.-C. Chang, and John R. Smith, ACM International Conf. on Multimedia (MM), pp.572-579, New York, October 2004.
54. Distance Function Design and Fusion for Sequence-data Mining, Y. Wu and E. Y. Chang, ACM International Conference on Information and Knowledge Management (CIKM), DC, November 2004.
55. Aligning Boundary in Kernel Space for Learning Imbalanced Dataset, G. Wu and E. Y. Chang, IEEE International Conference on Data Mining (ICDM), United Kingdom, November 2004.
56. Optimal Multimodal Fusion for Multimedia Data Analysis, Y. Wu, E. Y. Chang, K. Chang, and John Smith, ACM International Conference on Multimedia (MM), New York, October 2004.

57. Concept-dependent Multimodal Active Learning for Image Retrieval, K. Goh, E. Y. Chang, and W.-C. Lai, ACM International Conference on Multimedia (MM), New York, October 2004.
58. Adaptive Stream Resource Management Using Kalman Filters, A. Jain, E. Y. Chang, and Y.-F. Wang, ACM International Conference on Management of Data (SIGMOD), pp.11-22, Paris, June 2004.
59. Why One Example Is Not Enough for An Image Query, Thomas E. Bjoerge and Edward Y. Chang, IEEE International Conf. on Multimedia, Taipei, June 2004.
60. Distributed Video Data Fusion and Mining, E. Chang, Y.-F. Wang, and Volkan Rodoplu, SPIE Conf. on Defense and Security --- Sensors, and Command, Control, Communications, and Intelligence Technologies for Homeland Security and Homeland Defense (co-chaired by DARPA, Dod/Doj), Orlando, April 2004.
61. Using MEGA to Predict Molecular Bio-Activity, A. Qamra, Kingshy Goh, and E. Y. Chang, ACM International Conference on Research in Computational Molecular Biology, San Diego, March 2004.
62. Multi-camera Spatio-temporal Fusion and Biased Sequence-data Learning for Security Surveillance, G. Wu, Y. Wu, L. Jiao, Y.-F. Wang, and E. Chang, ACM International Conference on Multimedia, pp.528-538, Berkeley, November 2003.
63. Confidence-based Dynamic Ensemble for Image Annotation and Semantics Discovery, B. Li, K.-S. Goh, and E. Chang, ACM International Conference on Multimedia (**best student paper**), pp.195-206, Berkeley, November 2003.
64. Adaptive Feature-Space Conformal Transformation for Imbalanced-Data Learning, G. Wu and E. Chang, International Conference on Machine Learning (ICML), pp.816-823, DC, August 2003.
65. The Anatomy of A Multimodal Information Filter, Y.-L. Wu, K.-S. Goh, B. Li, H. You, E. Chang, ACM International Conference on Knowledge Discovery and Data Mining (KDD), pp.462-471, Washington DC, August 2003.
66. Invariant Feature Extraction and Biased Statistical Inference for Video Surveillance, Y. Wu, L. Jiao, G. Wu, E. Chang, and Y.-F. Wang, IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS), pp.284-289, Miami, July 2003.
67. Enhancing DPF for Near-replica Image Recognition, Y. Meng, E. Chang, and B. Li, IEEE International Conference on Computer Vision and Pattern Recognition (CVPR), pp.416-423, Madison, June 2003.
68. Design and Implementation of Semi-Preemptible IO, Z. Dimitrijevic, R. Rangaswami, and E. Y. Chang, USENIX Conf. on File and Storage Technologies (FAST), pp.145-158, SF, April 2003.
69. MEMS-based Disk Buffer for Streaming Media Servers, R. Rangaswami, Z. Dimitrijevic, E. Y. Chang, and K. E. Schausser, IEEE International Conference on Data Engineering (ICDE), pp.619-630, Bangalore, March 2003.
70. Image Copy Detection Using DPF, Y. Meng and E. Chang, IS&T/SPIE International Conference on Storage and Retrieval for Media Databases, pp.176-186, San Jose, January 2003.
71. A Framework for Detecting Hazardous Events, Y. Wu, G. Wu, and E. Chang, IS&T/SPIE International Conference on Storage and Retrieval for Media Databases, San Jose, January 2003.
72. DynDex: a Dynamic and Non-metric Space Indexer, K. Goh, B. Li, and E. Chang, ACM International Conference on Multimedia, pp.466-475, Juan Les Pin, France, December 2002 (14% accepted).
73. Virtual IO: Preemptible Disk Access (short paper), Z. Dimitrijevic, R. Rangaswami, and E. Chang, ACM International Conference on Multimedia, Juan Les Pin, December 2002.
74. An Anatomy of a Large-scale Image Search Engine, W.-C. Lai, E. Chang and K.-T. Cheng, IEEE MSE, Irvine, December 2002.
75. Hybrid Learning Schemes for Multimedia Information Retrieval, W.-C. Lai, E. Chang, and K.-T. Cheng, IEEE Pacific-Rim Conference on Multimedia, Hsing-Chu, December 2002.
76. MORF: A Distributed Multimodal Information Filtering System, Y.-L. Wu, E. Chang, et. al., IEEE Pacific-Rim Conf. on Multimedia, Hsing-Chu, December 2002.
77. An Architecture of a Web-based Collaborative Image Search Engine, W.-C. Lai, G. Sychay, and E. Chang, The 10th International Conference on Cooperative Information Systems, Irvine, October 2002.
78. DPF --- A Perceptual Distance Function for Image Retrieval, B. Li, E. Chang, C.-T. Wu, IEEE International Conference on Image Processing (ICIP), Rochester, September 2002.

79. BPMs vs. SVMs for Image Classification, G. Wu, E. Chang, and C.-S. Li, IEEE International Conference on Multimedia, Switzerland, August 2002.
80. Video Shot Transition Detection, Y. Wu, E. Chang, and B. Li, IEEE International Conference on Multimedia, Switzerland, August 2002.
81. The xTream Multimedia System, Z. Dimitrijevic, R. Rangaswami, and E. Chang, IEEE International Conference on Multimedia, Switzerland, August 2002.
82. Effective Image Annotation via Active Learning, G. Sychay, E. Chang, and K. Goh, IEEE International Conference on Multimedia, Switzerland, August 2002.
83. An Anatomy of a Large-scale Image Search Engine, W.-C. Lai, E. Chang, and K.-T. Cheng, WWW11, Hawaii, May 2002.
84. Supporting Subjective Image Queries without Seeding Requirements --- Proposing Test Queries for Benchathlon, E. Chang and K.-T. Cheng, IS&T/SPIE Internet Imaging III, pp.225-32, San Jose, Jan. 2002.
85. Spin Discriminant Analysis: Using a One Dimensional Classifier for High-Dimensional Classification Problems, H. You and E. Chang, IEEE CVPR, pp.968-975, Hawaii, December 2001.
86. Mining Image Features for Efficient Query Processing, B. Li, W.-C. Lai, E. Chang and K.-T. Cheng, IEEE Data Mining Conference, pp.353-360, San Jose, November 2001 (19% accepted).
87. SVM Binary Classifier Ensembles for Multi-Class Image Classification, K. Goh, E. Chang and K.-T. Cheng, ACM International Conference on Information and Knowledge Management (CIKM), pp.395-402, Atlanta, November 2001 (24% accepted).
88. Support Vector Machine Active Learning for Image Retrieval, S. Tong and E. Chang, ACM International Conference on Multimedia, pp.107-118, Ottawa, October 2001 (16% accepted).
89. Learning Image Query Concepts via Intelligent Sampling, B. Li, E. Chang, and C.-S. Li, IEEE International Conference on Multimedia, pp.1168-1171, Tokyo, August 2001.
90. NNEW - Nearest Neighbor Expansion by Weighting in Image Database Retrieval, H. You, E. Chang and B. Li, IEEE International Conference on Multimedia, pp.245-248, Tokyo, August 2001.
91. Data Placement for Multi-User Interactive Digital VCR, R. Rangaswami, E. Chang, C. Li and M. Chen, IEEE International Conf. on Multimedia, pp. 1128-1131, Tokyo, Aug., 2001.
92. Providing Scalable On-Demand Interactive Video Services by Means of Multicasting and Client Buffering. S.-H. Chan and E. Chang, IEEE International Conf. on Communications, Helsinki, June, 2001.
93. On Answering Queries in the Presence of Limited Access Patterns, C. Li and E. Chang, International Conf. on Database Theory (ICDT), London, January 2001.
94. Query Planning with Limited Source Capabilities, Chen Li and Edward Chang, IEEE International Conference on Data Engineering (ICDE), p. 401-412, San Diego, March 2000 (14% accepted).
95. PiDTV: A Client-Based Interactive DTV Architecture (short paper), Edward Chang, ACM International Conference on Multimedia, p.111-114, Orlando, November 1999.
96. Searching Near-Replicas of Images via Clustering, Edward Chang, Chen Li, James Z. Wang, Peter Mork, and Gio Wiederhold, IS&T/SPIE Symp. of Voice, Video, and Data Communications, p.281-92, Sept. 1999.
97. MEDIC: A Memory & Disk Cache for Multimedia Clients, Edward Chang and Hector Garcia-Molina, IEEE International Conf. on Multimedia Computing and Systems, p.493-99, Florence, June 1999.
98. On Managing Continuous Media Data, Edward Chang (Advisor: Hector Garcia Molina), Stanford University Ph.D. Dissertation, August 1999.
99. On Constructing Personalized Navigation Trees for Web Documents, K. Vu, W. Li, E. Chang, WWW Conference (poster), Toronto, May 1999.
100. RIME: A Replicated Image Detector for the World-Wide Web, Edward Chang, James Z. Wang, Chen Li, and Gio Wiederhold, IS&T/SPIE Symp. of Voice, Video, and Data Communications, November 1998.
101. 2D BubbleUp - Managing Parallel Disks for Media Servers, Edward Chang, Hector Garcia-Molina, and Chen Li, International Conference on Foundations of Data Organization (FODO), p221-230, Kobe, Japan, November 1998.

102. BubbleUp: Low Latency Fast-Scan for Media Servers, Edward Chang and Hector Garcia-Molina, ACM International Conference on Multimedia, p.87-98, Seattle, November 1997.
103. Effective Memory Use in a Media Server, Edward Chang and Hector Garcia-Molina, The 23rd Very Large Data Base (VLDB) Conference, p.496-505, Athens, Greece, August 1997.
104. Reducing Initial Latency in Media Servers, Edward Chang and Hector Garcia-Molina, IEEE Multimedia, vol.4, no.3, p.50-61, July 1997.
105. An Open and Extensible Event-Based Transaction Manager, E. Cheng, E. Chang and etc., USENIX Summer 1991, pp.49-58.

Refereed Workshop Publications

1. DeepQ Arrhythmia Database: A large-scale dataset for arrhythmia detector evaluation, Meng-Hsi Wu and Edward Chang, ACM MM MMHealth Workshop, Mountain View, October, 2017.
2. Inquire and Diagnose: Neural Symptom Checking Ensemble using Deep Reinforcement Learning, Kevin Tang, Haocheng Kao, Jason Chou, Edward Y. Chang, NIPS Workshop on Deep Reinforcement Learning, Barcelona, December 2016.
3. SpeedO - Parallelizing Stochastic Gradient Descent for Deep Convolutional Neural Network, Zhongyang Zheng, Wenrui Jiang, Gang Wu, Cyrus Choi, and Edward Y. Chang, NIPS Workshop on Learning Systems, Montreal, December 2015.
4. INS Assisted 360-Degree Panorama, Lun-Chen Chu, Andre Chen, Yuhsin Lin, Yu-Mei Chen, Scott Liao, and Edward Y. Chang, Workshop of Immersive Media Experience, in conjunction with ACM Multimedia, October 2015.
5. XINS – The Anatomy of An Indoor Positioning and Navigation Architecture, Y. Gao, Q. Yang, G. Li, E. Y. Chang, et al., ACM Workshop on Mobile Location-Based Service (MLBS), Beijing, September 2011.
6. INSPORAMA: INS-Aided Misalignment Correction in Feature-Based Panoramic Image Stitching, Yuan Gao, Chengu Wang, Edward Y. Chang, ICME Workshops 2012: 659-664.
7. Inertial Sensor Aided Alignment and Stitching for Panoramas on Mobile Phones, Q. Yang, H. Qiu, C. Wang, Y. Gao, and E. Y. Chang, ACM Workshop on Mobile Location-Based Service (MLBS), Beijing, September 2011.
8. Scalable Indexing for Perceptual Data (**invited**), A. Qamra, and E. Y. Chang, International Workshop on Multimedia Content Analysis and Mining, June 2007.
9. Speeding up SVMs, G. Wu, Z.-H. Zhang, and E. Y. Chang, Snowbird Machine Learning Workshop, April 2006.
10. Unified Learning Paradigm for Web-Scale Mining, E. Y. Chang and M. Lyu, Snowbird Machine Learning Workshop, April 2006.
11. EXTENT for Photo Annotation, E. Y. Chang, Workshop on Computer Vision Meets Databases (CVDB) in cooperation with ACM International Conference on Management of Data (SIGMOD), Baltimore, June 2005.
12. Using Pivots to Index for Support Vector Machine Queries, A. Qamra and E. Y. Chang, Workshop on Computer Vision Meets Databases (CVDB) in cooperation with ACM International Conference on Management of Data (SIGMOD), Baltimore, June 2005.
13. Adaptive Sampling for Sensor Networks, A. Jain and E. Y. Chang, International Workshop on Data Management for Sensor Networks in conjunction with VLDB, Toronto, August 2004.
14. On Scalability of Active Learning for Formulating Query Concepts, W.-C. Lai, Kingshy Goh, and E. Y. Chang, Workshop on Computer Vision Meets Databases (CVDB) in corporation with ACM International Conference on Management of Data (SIGMOD), pp.11-18, Paris, June 2004.
15. Class-boundary Alignment for Imbalanced Dataset Learning, G. Wu and E. Chang, International Conference on Machine Learning (ICML) Workshop on Learning from Imbalanced Datasets, pp.49-56, Washington DC, August 2003.
16. Support Vector Machine Pair-wise Classifiers with Error Reduction for Image Classification, K. Goh, E. Chang and K.-T. Cheng, ACM International Conference on Multimedia (MIR Workshop), Ottawa, October 2001.

17. Towards Perception-Based Image Retrieval, Edward Chang, Beita Li and Chen Li, IEEE Workshop on Content-based Access of Image and Video Libraries, p.101-105, South Carolina, June, 2000.
18. An Image Coding and Reconstruction Scheme for Mobile Computing, Edward Chang, The 5th IDMS (Springer-Verlag LNCS #1483), p.137-148, Oslo, Norway, September 1998. (Best papers selected for journal submission)
19. Cost-Based Media Server Design, Edward Chang and Hector Garcia-Molina, The 8th Research Issues in Data Engineering, p.76-83, Orlando, February 1998 (Best papers selected for journal submission).
20. Minimizing Initial Latency in a Multimedia Storage System, Edward Chang and Hector Garcia-Molina, International Workshop on Multimedia Database Systems, August 1996.

Demonstrations

1. INS Assisted 360-Degree Panorama, Lun-Chen Chu, Andre Chen, Yuhsin Lin, Yu-Mei Chen, Scott Liao, and Edward Y. Chang, ACM International Conf. on Multimedia, October 2015. (**Best technical demo award**)
2. Fotowiki - Distributed Map Enhancement Service, Wen-Yen Chen and E. Y. Chang, ACM International Conference on Multimedia (MM), October 2006.
3. Fotofiti, B. Lee, W.-Y. Chen, and E. Y. Chang, ACM International Conference on Multimedia (MM), October 2006. A Multimodal Image Database System, E. Chang, B. Li, W.-C. Lai, C. Chang, K.-T. Cheng, and M. Crandell (with VIMA Technologies), IEEE International Conference on Computer Vision and Pattern Recognition (CVPR), Madison, June 2003.
4. PBIR-MM: Multimodal Image Retrieval and Annotation, W.-C. Lai, C. Chang, E. Chang, K.-T. Cheng, and M. Crandell, ACM International Conf. on Multimedia, Juan Les Pin, Dec. 2002.
5. PBIR: A System that Learns Subjective Image Query Concepts, E. Chang, K.-T. Cheng, W.-C. Lai, C. Wu, C. Chang and Y. Wu, ACM International Conference on Multimedia, pp.611-614, Ottawa, October 2001..
6. PBIR - Perception-Based Image Retrieval, E. Chang, K.-T. Cheng and L. Chang, ACM SIGMOD, Santa Barbara, May 2001.
7. PowerBookmarks: A System for Personalizable Web Information Organization, Sharing, and Management, Quoc Vu, Wen-Syan Li and Edward Chang, ACM SIGMOD, p.565-67, Philadelphia, June 1999.

Grants

1. 2013-2016: Mobile and Wearable Devices for Healthcare (PI, Taiwan Ministry of Economy; US\$6,000,000)
2. 2006-2008: Scalable, Multimodal Algorithms for Multimedia Information Retrieval (PI, NSF IIS-0535085, Google; US\$300,000)
3. 2004-2006: VSN: Video Sensor Networks (Proximex, Intel; US\$300,000)
4. 2003-2005: ISE: Interactive Streaming Everywhere (Participant, NSF EIA-0080134; 2000-2003; SONY, UC DiMI; US\$500,000)
5. 2004-2006: SIEC: Statistical Inferences under Extreme Constraints (PI, NSF ITR IIS-0219885, ARDA, NEC, IBM; US\$300,000)
6. 2002-2005: Intelligent Sampling for Learning Image Query Concepts (PI, NSF Career IIS-0133802; US\$300,000)

Ph.D. Thesis Advising and Collaborators

- **Google Intern Advisees** (Full-Time Interns at Google later admitted to top PhD programs): Haoyuan Li (PhD CS, UC Berkeley), Young Gao (PhD CS, MIT), Hao Su (PhD CS, Stanford), Mike Stanton (PhD CS, CMU), Junchen Jiang (PhD CS, CMU), Yuqian Li (PHD CS, Duke).

- **UCSB Graduated Ph.D. Advisees:** Huaxin Yu (2002, Nelson)), Beita Li (2003, ASK), Zoran Dimitrijevic (2004, Google), Raju Rangaswami (2004, Florida U), Kingshy Goh (2004), Yi Wu (2005, Intel Research), Gang Wu (2006, Microsoft), Ankur Jain (2006), Navneet Panda (2006, Google), Ben Lee (2006, Google), Arun Qamra (2007, Google), Wen-Yen Chen (2009, Yahoo), Xiance Si (2010, Google), Zhiyuan Liu (2011, Tsinghua U).
- **MIT Advisees** (MIT MS students visited Google for a year): Jon Chu (2007-09), Rong Hu (2008-09), Brian Wong (2008-10), Alice Li (2008-10).