

Rogério Schmidt Feris

March 2021

MIT-IBM Watson AI Lab, IBM Research
75 Binney St.
Cambridge, MA 02142, USA

Tel: 914-255-0218
rsferis@us.ibm.com
<http://www.rogerioferis.org>

Education

Ph.D. degree in Computer Science (2002-2006)

University of California, Santa Barbara, USA
(Advised by Prof. Matthew Turk)

M.S. degree in Computer Science (1999-2001)

University of Sao Paulo, Brazil
(Advised by Prof. Roberto Cesar)

B.S. degree in Computer Engineering (1994-1998)

University of Rio Grande (FURG), Brazil

Professional Experience

- **Principal Scientist & Manager, MIT-IBM Watson AI Lab, MA (2020-present)**
 - I am currently leading a team of researchers from IBM and MIT as a principal investigator for DARPA Learning with Less Labels (LwLL), together with Prof. Josh Tenenbaum. I also lead the IBM-Purdue team as a principal investigator for IARPA DIVA (with Prof. Jeff Siskind) and the MIT-IBM Moments in Time project (with Prof. Aude Oliva). We have achieved strong results in the evaluation benchmarks for both LwLL and DIVA, ranking 1st or 2nd for all tasks among all other performers and external institutions.
- **Research Manager, IBM T. J. Watson Research Center, NY (2016-2020)**
 - I led the research team that created a system called Cognitive Highlights, which was used to generate the official highlights of the US Open/Wimbledon/Masters tournaments. The highlights were seen by millions of people! This work was covered by the New York Times and many other media outlets. It was also integrated as a product offering of Watson Media. We received a prestigious Corporate award, which recognizes the impact of this work.
 - I served as the coordinator of the Joint Program between Research and Watson Visual Recognition, driving efforts towards strong research and product deliverables for the Watson Vision Services. I was also a global leader for the multimodal perception area at IBM, coordinating efforts across labs worldwide and leading computer vision / AI challenges.
- **Research Scientist, IBM T. J. Watson Research Center, NY (2006-2016)**
 - I was one of the key inventors of the IBM Intelligent Video Analytics (IVA), a leading-edge camera-based system with automatic event detection and search capabilities. My work on video analysis for safety and security was published in high-

profile conferences and featured at ABC news and CBS 60 minutes. Among other awards, I was named IBM Master Inventor and received an IBM Outstanding Innovation Achievement Award.

- **Affiliate Associate Professor, University of Washington, WA (2008-2016)**
 - I have co-advised PhD students at University of Washington, in close collaboration with Prof. Ming-Ting Sun. I co-advised the PhD dissertation of Haowei Liu and Jun Xie (both working at Google now).
- **Adjunct Associate Professor, Columbia University, NY (2008, 2013, 2014)**
 - In 2013 and 2014 I co-taught advanced graduate courses on “Visual Recognition and Search” (Electrical Engineering and Computer Science Departments). In 2008 I co-taught a graduate course on “Automatic Video Surveillance” (Computer Science Department).
- **Senior Software Engineer, IBM Global Technology Services, NY (2008)**
 - In addition to working on core research, I had a one-year assignment to IBM Global Technology Services as a senior software engineer, with the goal of helping the commercialization of the IBM Intelligent Video Analytics (IVA) system. My code has been running in various components of the product.

Research Interests

- **Broad Research Interests:**
 - Computer Vision, Machine Learning, Multimedia, Artificial Intelligence
- **Specific Research Interests:**
 - Learning with Limited Labels, Transfer Learning, Multi-Task Learning
 - Multimodal Video and Image Analysis
 - Deep Learning, Dynamic Neural Networks, Efficient AI
 - Visual Analysis of People; Computer Vision for Fashion
 - Real-world Applications of Computer Vision

Selected Awards and Distinctions

- IBM Corporate Award for AI Video Enrichment and Editing (Auto-curation of Sports Highlights), 2019. This is a prestigious company-wide award which recognizes breakthrough achievements that have led to significant market and industry success for IBM.
- 1st place, TRECVID ActEV: Activities in Extended Video, 2018.
 - Activity Detection (AD) Task – 1st Place out of 30 entries.
 - Activity and Object Detection (AOD) Task – 1st Place out of 15 entries
- Best Poster Award, Greater New York Area Workshop on Multimedia and Vision, 2013. For the paper "Designing category-level attributes for discriminative visual recognition", with Felix Yu, Liangliang Cao, John Smith, and Shih-Fu Chang.

- Invited Speaker at the National Academy of Engineering JAF0E meeting, which gathered sixty of the most promising engineers under the age of 45 from Japan and the United States, 2012.
- TRECVID Surveillance Event Detection Competition (IBM-CMU-NUS team). I was part of the team that ranked 1st in 4 out of 7 event detection tasks, 2012.
- Named IBM Master Inventor, 2011. This title is reserved for leading inventors within IBM. Each year, only 1% of the IBM Research technical population receives this award.
- IBM Outstanding Innovation Achievement Award. IBM Research award received for contributions in physical security and video analytics, 2011.
- Far-Reaching Research Award – competitive grant from IBM Research to pursue research towards building the next-generation camera networks (with Ankur Datta, Lisa Brown, and Rick Kjeldsen), 2010.
- North American Frost & Sullivan Award for Video Surveillance Software Company of the Year, 2008. Prize given to IBM in recognition of our work in video surveillance.
- IBM Emerging Leader in Multimedia, July 2005. I was among eight students selected from top universities by IBM.
- Best Computer Science MSc Thesis in Brazil - 2nd Prize (Nationwide Competition). Award received at the Congress of Brazilian Computer Society, 2002.
- IBM Research Accomplishment Awards
 - Dynamic Neural Networks for Efficient AI, 2020
 - Contributions to Science in the Learning with Less Labels (LwLL) Domain – Data Limited Learning for Vision and Beyond, 2020
 - Efficient and Robust Object Detection, 2018
 - CoreML, Customized Learning and De-Biasing for Watson Visual Recognition, 2018
 - Research Contributions to Watson Visual Recognition Services, 2017
 - Contributions to Safety & Physical Security through Video Analytics (O-level, Outstanding), 2010

Publications

[140+ papers, h-index 47 as of 2020]

Books

1. R. Feris, C. Lampert, and D. Parikh (Eds.). Visual Attributes. Advances in Computer Vision and Pattern Recognition, Springer, 2016.

Peer-Reviewed Conference Publications

1. G. Bukchin, E. Schwartz, K. Saenko, O. Shahar, R. Feris, R. Giryes, and L. Karlinsky. Fine-grained Angular Contrastive Learning with Coarse Labels. Conference on Computer Vision and Pattern Recognition (**CVPR 2021, Oral**).

2. H. Wu, Y. Gao, X. Guo, Z. Al-Halah, S. Rennie, K. Grauman, and R. Feris. Fashion IQ: A New Dataset Towards Retrieving Images by Natural Language Feedback. Conference on Computer Vision and Pattern Recognition (**CVPR 2021**).
3. S. Whitehead, H. Wu, H. Ji, R. Feris, and K. Saenko. Separating Skills and Concepts for Novel Visual Question Answering. Conference on Computer Vision and Pattern Recognition (**CVPR 2021**).
4. M. Monfort, S. Jin, D. Harwath, R. Feris, J. Glass, and A. Oliva. Spoken Moments: A Large Scale Dataset of Audio Descriptions of Dynamic Events in Video. Conference on Computer Vision and Pattern Recognition (**CVPR 2021**).
5. A. Singh, O. Chakraborty, A. Varshney, R. Panda, R. Feris, K. Saenko, and A. Das. Semi-Supervised Action Recognition with Temporal Contrastive Learning. Conference on Computer Vision and Pattern Recognition (**CVPR 2021**).
6. C. Chen, R. Panda, K. Ramakrishnan, R. Feris, J. Cohn, A. Oliva, and Q. Fan. Deep Analysis of CNN-based Spatio-temporal Representations for Action Recognition. Conference on Computer Vision and Pattern Recognition (**CVPR 2021**).
7. B. Pan, R. Panda, C. Fosco, C. Lin, A. Andonian, Y. Meng, K. Saenko, A. Oliva, and R. Feris. VA-RED²: Video Adaptive Redundancy Reduction. International Conference on Learning Representations (**ICLR 2021**).
8. Y. Meng, R. Panda, C. Lin, P. Sattigeri, L. Karlinsky, K. Saenko, A. Oliva, and R. Feris. AdaFuse: Adaptive Temporal Fusion Network for Efficient Action Recognition. International Conference on Learning Representations (**ICLR 2021**).
9. L. Karlinsky, J. Shtok, A. Alfassy, M. Lichtenstein, S. Harary, E. Schwartz, S. Doveh, P. Sattigeri, R. Feris, A. Bronstein, and R. Giryes. StarNet: towards Weakly Supervised Few-Shot Object Detection. AAAI Conference on Artificial Intelligence (**AAAI 2021**).
10. R. Panda, M. Merler, M. Jaiswal, H. Wu, K. Ramakrishnan, U. Finkler, C. Chen, M. Cho, R. Feris, D. Kung, and B. Bhattacharjee. NASTransfer: Analyzing Architecture Transferability in Large Scale Neural Architecture Search. AAAI Conference on Artificial Intelligence (**AAAI 2021**).
11. U. Finkler, M. Merler, R. Panda, M. Jaiswal, H. Wu, K. Ramakrishnan, C. Chen, M. Cho, D. Kung, R. Feris, and B. Bhattacharjee. Large Scale Neural Architecture Search with Polyharmonic Splines. AAAI Workshop on Meta-Learning for Computer Vision, 2021.
12. X. Sun, R. Panda, R. Feris, and K. Saenko. AdaShare: Learning What to Share for Efficient Deep Multi-Task Learning. Conference on Neural Information Processing Systems (**NeurIPS 2020**).
13. Y. Guo, N. Codella, L. Karlinsky, J. Codella, J. Smith, K. Saenko, T. Rosing, and R. Feris. A Broader Study of Cross-Domain Few-Shot Learning. European Conference on Computer Vision (**ECCV 2020**).
14. Y. Meng, C. Lin, R. Panda, P. Sattigeri, L. Karlinsky, K. Saenko, A. Oliva, and R. Feris. AR-Net: Adaptive Frame Resolution for Efficient Action Recognition. European Conference on Computer Vision (**ECCV 2020**).

15. Z. Tang, Y. Gao, P. Sattigeri, L. Karlinsky, R. Feris, and D. Metaxas. OnlineAugment: Online Data Augmentation with Less Domain Knowledge. European Conference on Computer Vision (**ECCV 2020**).
16. M. Lichtenstein, P. Sattigeri, R. Feris, R. Giryes, and L. Karlinsky. TAFSSL: Task-Adaptive Feature Sub-Space Learning for few-shot classification. European Conference on Computer Vision (**ECCV 2020**).
17. A. Andonian, C. Fosco, M. Monfort, A. Lee, R. Feris, C. Vondrick, and A. Oliva. We Have So Much In Common: Modeling Semantic Relational Set Abstractions in Videos. European Conference on Computer Vision (**ECCV 2020**).
18. A. Sahoo, A. Singh, R. Panda, R. Feris, and A. Das. Mitigating Dataset Imbalance via Joint Generation and Classification. ECCV Workshop on Imbalance Problems in Computer Vision, 2020.
19. C. Lin, Y. Hung, R. Feris, and L. He. Video Instance Segmentation Tracking. Conference on Computer Vision and Pattern Recognition (**CVPR 2020**).
20. Z. Wang, M. Yu, Y. Wei, R. Feris, J. Xiong, W. Hwu, T. Huang, and H. Shi. Differential Treatment for Stuff and Things: A Simple Unsupervised Domain Adaptation Method for Semantic Segmentation. Conference on Computer Vision and Pattern Recognition (**CVPR 2020**).
21. K. Ramakrishnan, R. Panda, Q. Fan, J. Henning, A. Oliva, and R. Feris. Relationship Matters: Relation Guided Knowledge Transfer for Incremental Learning of Object Detectors. CVPR Workshop on Continual Learning in Computer Vision, 2020.
22. M. Khoi-Nguyen, D. Joshi, R. Yeh, J. Xiong, R. Feris, and M. Do. Learning Motion in Feature Space: Locally- Consistent Deformable Convolution Networks for Fine Grained Action Detection. International Conference on Computer Vision (**ICCV 2019, Oral**), Seoul, Korea, 2019.
23. A. Alfassy, L. Karlinsky, A. Aides, J. Shtok, S. Harary, R. Feris, R. Giryes, and A. Bronstein. LaSO: Label-Set Operations Networks for Multi-label Few-shot Learning. Conference on Computer Vision and Pattern Recognition (**CVPR 2019, Oral**), Long Beach, California, 2019.
24. Y. Guo, H. Shi, A. Kumar, K. Grauman, T. Rosing, and R. Feris. SpotTune: Transfer Learning through Adaptive Fine-tuning. Conference on Computer Vision and Pattern Recognition (**CVPR 2019**), Long Beach, California, 2019.
25. L. Karlinsky, J. Shtok, S. Harary, E. Schwartz, A. Aides, R. Feris, R. Giryes, and A. Bronstein. RepMet: Representative-based Metric Learning for Classification and One-shot Object Detection. Conference on Computer Vision and Pattern Recognition (**CVPR 2019**), Long Beach, California, 2019.
26. C. Chen, Q. Fan, N. Mallinar, T. Sercu, and R. Feris. Big-Little Net: An Efficient Multi-Scale Feature Representation for Visual and Speech Recognition. International Conference on Learning Representations (**ICLR 2019**), New Orleans, Louisiana, 2019.
27. A. Boggust, K. Audhkhasi, D. Joshi, D. Harwath, S. Thomas, R. Feris, D. Gutfreund, Y. Zhang, A. Torralba, M. Picheny, and James Glass. Grounding Spoken Words in Unlabeled Video. CVPR Workshop on Sight and Sound, 2019.

28. K. Ramakrishnan, M. Monfort, B. McNamara, A. Lascelles, D. Gutfreund, R. Feris, and A. Oliva. Identifying Interpretable Action Concepts in Deep Networks. CVPR Workshop on Explainable AI, 2019.
29. M. Jaiswal et al. Video-Text Compliance: Activity Verification Based on Natural Language Instructions. ICCV Workshop on Large Scale Holistic Video Understanding, 2019.
30. Z. Shen, H. Shi, J. Yu, H. Phan, R. Feris, L. Cao, D. Liu, X. Wang, T. Huang, M. Savvides. Learning Object Detection from Scratch via Gated Feature Reuse. British Machine Vision Conference (BMVC), 2019.
31. E. Schwartz*, L. Karlinsky*, J. Shtok, S. Harary, M. Marder, A. Kumar, R. Feris, R. Giryes, and A. Bronstein. Delta-Encoder: an Effective Sample Synthesis Method for Few-shot Object Recognition. Neural Information Processing Systems (**NeurIPS 2018, Spotlight**), Montreal, Canada, 2018. (* equal contribution)
32. X. Guo*, H. Wu*, Y. Cheng, S. Rennie, G. Tesauro, and R. Feris. Dialog-based Interactive Image Retrieval. Neural Information Processing Systems (**NeurIPS 2018**), Montreal, Canada, 2018. (*equal contribution)
33. A. Kumar, P. Sattigeri, K. Wadhawan, L. Karlinsky, R. Feris, W. T. Freeman, and G. Wornell. Co-regularized Alignment for Unsupervised Domain Adaptation. Neural Information Processing Systems (**NeurIPS 2018**), Montreal, Canada, 2018.
34. R. Gao, R. Feris and K. Grauman. Learning to Separate Object Sounds by Watching Unlabeled Video. European Conference on Computer Vision (**ECCV 2018, Oral**), Munich, Germany, 2018.
35. B. Cheng*, Y. Wei*, H. She, R. Feris, J. Xiong, and T. Huang. Revisiting RCNN: On Awakening the Classification Power of Faster RCNN. European Conference on Computer Vision (**ECCV 2018**), Munich, Germany, 2018. (* equal contribution)
36. Z. Wu*, T. Nagarajan*, A. Kumar, S. Rennie, L. Davis, K. Grauman, and R. Feris. BlockDrop: Dynamic Inference Paths in Residual Networks. IEEE Conference on Computer Vision and Pattern Recognition (**CVPR 2018, Spotlight**), Salt Lake City, Utah, June 2018. (* equal contribution)
37. X. Peng, Z. Tang, F. Yang, R. Feris, and D. Metaxas. Jointly optimize data augmentation and network training: Adversarial data augmentation in human pose estimation. IEEE Conference on Computer Vision and Pattern Recognition (**CVPR 2018**), Salt Lake City, Utah, June 2018.
38. M. Merler, D. Joshi, K. Mac, Q. Nguyen, S. Hammer, J. Kent, J. Xiong, M. Do, J. Smith, and R. Feris. The Excitement of Sports: Automatic Highlights using Audio-Visual Cues. CVPR Workshop on Sight and Sound, 2018.
39. N. Codella, C. Lin, A. Halpern, M. Hind, R. Feris, and J. Smith. Collaborative Human-AI (CHAI): Evidence-Based Interpretable Melanoma Classification in Dermoscopic Images. MICCAI Workshop on Interpretability of Machine Intelligence in Medical Image Computing, 2018.
40. N. Codella, D. Anderson, T. Philips, A. Porto, K. Massey, J. Snowdon, R. Feris, and J. Smith. Segmentation of both Diseased and Healthy Skin from Clinical Photographs in a Primary Care Setting. International Engineering in Medicine and Biology Conference, 2018.

41. M. Beigi, L. M Brown, Q. Fan, J. Henning, C. Lin, H. Shi, C. Shu, and R. Feris. Object-Centric Spatio-Temporal Activity Detection and Recognition. NIST TRECVID Workshop, 2018.
42. Y. Lu, A. Kumar, S. Zhai, Y. Cheng, T. Javidi, and R. Feris. Fully Adaptive Feature Sharing in Multi-Task Networks with Applications in Person Attribute Classification. IEEE Conference on Computer Vision and Pattern Recognition (**CVPR 2017, Spotlight**), Honolulu, Hawaii, July 2017.
43. S. Zhai, H. Wu, A. Kumar, Y. Cheng, Y. Lu, Z. Zhang, and R. Feris. S3Pool: Pooling with Stochastic Spatial Sampling. IEEE Conference on Computer Vision and Pattern Recognition (**CVPR 2017**), Honolulu, Hawaii, July 2017.
44. M. Merler, D. Joshi, Q. Nguyen, S. Hammer, J. Kent, J. R Smith, and R. Feris. Automatic Curation of Golf Highlights using Multimodal Excitement Features. Workshop of Computer Vision in Sports (in conjunction with CVPR), 2017.
45. Z. Cai, Q. Fan, R. Feris, and N. Vasconcelos. A Unified Multi Scale Deep Convolutional Neural Network for Fast Object Detection. European Conference on Computer Vision (**ECCV 2016**), Amsterdam, Netherlands, 2016.
46. X. Peng, R. Feris, X. Wang, and D. Metaxas. A Recurrent Encoder-Decoder Network for Sequential Face Alignment. European Conference on Computer Vision (**ECCV 2016, Oral**), Amsterdam, Netherlands, 2016.
47. J. Wang, Y. Cheng, and R. Feris. Walk and Learn: Facial Attribute Representation Learning from Egocentric Video and Contextual Data. IEEE Conference on Computer Vision and Pattern Recognition (**CVPR 2016, Oral**), Las Vegas, Nevada, June 2016.
48. Y. Cheng, F. Yu, R. Feris, S. Kumar, A. Choudhary, and S. F. Chang. An Exploration of Parameter Redundancy in Deep Networks with Circulant Projections. IEEE International Conference on Computer Vision (**ICCV 2015**), Santiago, Chile, December 2015.
49. J. Huang, R. Feris, Q. Chen, and S. Yan. Cross-domain Image Retrieval with a Dual Attribute-aware Ranking Network. IEEE International Conference on Computer Vision (**ICCV 2015**), Santiago, Chile, December 2015.
50. Q. Chen, J. Huang, R. Feris, L. Brown, J. Dong, and S. Yan. Deep Domain Adaptation for Describing People Based on Fine-Grained Clothing Attributes. IEEE Conference on Computer Vision and Pattern Recognition (**CVPR 2015**), Boston, Massachusetts, June 2015.
51. J. Bowler, R. Feris, L. Cao, J. Wang, and M. Zhou. Automated Axon Segmentation from Highly Noisy Microscopic Videos. Winter Conference on Applications of Computer Vision (WACV 2015), Kona, Hawaii, 2015.
52. R. Feris, R. Bobbitt, and S. Pankanti. Efficient 24/7 Object Detection in Surveillance Videos. IEEE International Conference on Advanced Video and Signal-Based Surveillance (AVSS 2015), Germany, August 2015.
53. Y. Cheng, L. Brown, Q. Fan, R. Feris, S. Pankanti, and T. Zhang. RiskWheel: Interactive Visual Analytics for Surveillance Event Detection. IEEE International Conference on Multimedia and Expo (ICME 2014, Oral), Chengdu, China, 2014.

54. J. Xie, L. Chou, R. Feris, and M.T. Sun. Single Depth Image Super resolution and Denoising via Coupled Dictionary Learning with Local Constraints and Shock Filtering. IEEE International Conference on Multimedia and Expo (ICME 2014, Oral), 2014.
55. J. Xie, R. Feris, and M.T. Sun. Edge Guided Single Depth Image Super Resolution. IEEE International Conference on Image Processing (ICIP 2014), 2014.
56. Y. Cui, Y. Xiang, K. Rong, L. Cao, and R. Feris. A Spatial-Color Layout Feature for Representing Galaxy Images. IEEE Winter Conference on Applications of Computer Vision (WACV 2014), Steamboat Springs, Colorado, 2014.
57. L. Brown, R. Feris, and S. Pankanti. Temporal Non-Maximum Suppression for Pedestrian Detection Using Scene Context. International Conference on Pattern Recognition (ICPR 2014), Stockholm, Sweden, 2014.
58. R. Feris, L. Brown, S. Pankanti, and M.T. Sun. Appearance-based Object Detection under Varying Environmental Conditions. International Conference on Pattern Recognition (ICPR 2014, Oral), Stockholm, Sweden, 2014.
59. A. Mattos and R. Feris. Fusing Well-Crafted Feature Descriptors for Efficient Fine-Grained Classification. IEEE International Conference on Image Processing (ICIP 2014), Paris, France 2014.
60. A. Mattos, R. Herrmann, K. Shigeno, and R. Feris. A Mission-Oriented Citizen Science Platform for Efficient Flower Classification Based on Combination of Feature Descriptors. ICMR Workshop on Environmental Multimedia Retrieval, Glasgow, UK, 2014.
61. A. Mattos, R. Herrmann, K. Shigeno, and R. Feris. Flower Classification for a Citizen Science Mobile App. International Conference on Multimedia Retrieval (ICMR 2014), Glasgow, UK, 2014.
62. R. Feris, R. Bobbit, L. Brown, and S. Pankanti. Attribute-based People Search: Lessons Learnt from a Practical Surveillance System. ACM International Conference on Multimedia Retrieval (ICMR 2014), Oral Presentation, Glasgow, UK, 2014.
63. K. Scherbaum, R. Feris, J. Petterson, V. Blanz, and H. Seidel. Fast Face Detector Training Using Tailored Views. IEEE International Conference on Computer Vision (**ICCV 2013**), Sydney, Australia, December 2013.
64. Q. Chen, Z. Song, R. Feris, A. Datta, L. Cao, Z. Huang, and S. Yan. Efficient Maximum Appearance Search for Large-Scale Object Detection. IEEE Conference on Computer Vision and Pattern Recognition (**CVPR 2013**), Portland, Oregon, June 2013.
65. F. Yu, L. Cao, R. Feris, J. Smith, and S. Chang. Designing Category-Level Attributes for Discriminative Visual Recognition. IEEE Conference on Computer Vision and Pattern Recognition (**CVPR 2013**), Portland, Oregon, June 2013.
66. Q. Chen et al. Spatio-Temporal Fisher Vector Coding for Surveillance Event Detection. ACM International Conference on Multimedia (ACM MM 2013), Barcelona, Spain, 2013.
67. J. Xie, J. Hsu, R. Feris, and M. Sun. Fine Registration of 3D Point Clouds with ICP Using an RGB-D Camera. IEEE International Symposium on Circuits and Systems (ISCAS 2013), Beijing, China, 2013.

68. X. Wang, L. Cao, R. Feris, A. Datta, and T. Xan. Hierarchical Feature Pooling with Structure Learning: A new method for Pedestrian Detection. CVPR Workshop on Structured Prediction: Tractability, Learning, and Inference, Portland, Oregon, 2013.
69. J. Leandro, R. Cesar, and R. Feris. Shape Analysis using the Spectral Graph Wavelet Transform. IEEE eScience Conference, Beijing, China 2013.
70. R. Feris, A. Datta, M. T. Sun, and S. Pankanti. Boosting Object Detection Performance in Crowded Surveillance Videos. Winter Conference on Applications of Computer Vision (WACV 2013), Florida, USA, 2013.
71. F. Rashed, B. Siddiquie, R. Feris, and L. Davis. Domain Adaptive Object Detection. Winter Conference on Applications of Computer Vision (WACV 2013), Florida, USA, 2013.
72. Y. Cai et al. CMU-IBM-NUS@TRECVID 2012: Surveillance Event Detection. NIST Technical Report, 2012. First place in the retrospective surveillance event detection task.
73. R. Feris, B. Siddiquie, and S. Pankanti. Learning Detectors from Large Datasets for Object Retrieval in Video Surveillance. International Conference on Multimedia and Expo (ICME 2012), Melbourne, Australia, 2012.
74. B. Siddiquie, R. Feris, A. Datta, and L. Davis. Unsupervised Model Selection for View-Invariant Object Detection in Surveillance Environments. International Conference on Pattern Recognition (ICPR 2012, Oral), Tsukuba City, Japan, 2012.
75. A. Datta, L. Brown, R. Feris, and S. Pankanti. Appearance Modeling for Person Re-Identification using Weighted Brightness Transfer Functions. International Conference on Pattern Recognition (ICPR 2012), Tsukuba City, Japan, 2012.
76. B. Siddiquie, R. Feris, and L. Davis. Image Ranking and Retrieval Based on Multi-Attribute Queries. IEEE Conference on Computer Vision and Pattern Recognition (**CVPR 2011, Oral Presentation**), Colorado Springs, USA, 2011.
77. R. Feris, B. Siddiquie, Y. Zhai, J. Petterson, L. Brown, and S. Pankanti. Attribute-based Vehicle Search in Crowded Surveillance Videos. ACM International Conference on Multimedia Retrieval (ICMR 2011, Oral Presentation), Trento, Italy, 2011.
78. A. Datta, R. Feris, and D. Vaquero. Hierarchical Ranking of Facial Attributes. IEEE International Conference on Automatic Face and Gesture Recognition (FG 2011), Santa Barbara, California, 2011.
79. R. Feris, J. Petterson, B. Siddiquie, L. Brown, and S. Pankanti. Large-Scale Vehicle Detection in Challenging Urban Surveillance Environments. Winter Conference on Applications of Computer Vision (WACV 2011), Kona, Hawaii, 2011.
80. H. Liu, R. Feris, V. Krueger, and M.T. Sun. Unsupervised Action Classification Using Space-Time Link Analysis. IEEE International Symposium on Circuits and Systems (ISCAS 2010), Paris, France, 2010.
81. L. Chen, J. McAuley, R. Feris, T. Caetano, and M. Turk. Shape Classification Through Structured Learning of Matching Measures. IEEE Conference on Computer Vision and Pattern Recognition (**CVPR 2009**), Miami, Florida, June 2009.

82. D. Vaquero, R. Raskar, R. Feris, and M. Turk. A Projector-Camera Setup for Geometry-Invariant Frequency Demultiplexing. IEEE Conference on Computer Vision and Pattern Recognition (**CVPR 2009**), Miami, Florida, June 2009.
83. D. Vaquero, R. Feris, L. Brown, and A. Hampapur. Attribute-based people search in surveillance environments. Winter Conference on Applications of Computer Vision (WACV 2009, Oral Presentation), Snowbird, Utah, December 2009.
84. A. Hampapur, R. Bobbitt, L. Brown, M. Desimone, R. Feris, R. Kjeldsen, M. Lu, C. Mercier, C. Milite, S. Russo, C. Shu, Y. Zhai. Video Analytics in Urban Environments. IEEE International Conference on Advanced Video and Signal Based Surveillance, Genova, Italy, 2009.
85. R. Feris, Y. Tian, Y. Zhai, and A. Hampapur. Facial Image Analysis Using Local Feature Adaptation Prior to Learning. IEEE International Conference on Automatic Face and Gesture Recognition, Amsterdam, Netherlands, September 2008.
86. Y. Tian, R. Feris, and A. Hampapur. Real-Time Detection of Abandoned and Removed Objects in Complex Environments. IEEE International Workshop on Visual Surveillance (in conjunction with ECCV 2008), Marseille, France, 2008.
87. D. Vaquero, R. Feris, M. Turk, and R. Raskar. Characterizing the Shadow Space of Camera-Light Pairs. IEEE Conference on Computer Vision and Pattern Recognition (**CVPR 2008**), Anchorage, Alaska, June 2008.
88. L. Chen, R. Feris, Y. Zhai, L. Brown, and A. Hampapur. An Integrated System for Moving Object Classification in Surveillance Videos. IEEE International Conference on Advanced Video and Signal-Based Surveillance (AVSS 2008), Santa Fe, New Mexico, September 2008.
89. L. Chen, R. Feris, and M. Turk. Efficient Partial Shape Matching Using the Smith-Waterman Algorithm. CVPR Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment, Anchorage, Alaska, June 2008.
90. A. Hampapur, L. Brown, R. Feris, A. Senior, C. Shu, Y. Tian, Y. Zhai, and M. Lu. Searching Surveillance Video. IEEE International Conference on Advanced Video and Signal-Based Surveillance (AVSS 2007), London, UK, September 2007.
91. A. Senior, L. Brown, A. Hampapur, C. Shu, Y. Zhai, R. Feris, Y. Tian, S. Borger, and C. Carlson. Video Analytics for Retail: the IBM Smart Surveillance Retail Solution. IEEE International Conference on Advanced Video and Signal-Based Surveillance (AVSS 2007), London, UK, September 2007.
92. R. Feris, Y. Tian, and A. Hampapur. Capturing People in Surveillance Video. IEEE International Workshop on Visual Surveillance, 2007.
93. R. Feris, R. Raskar, and M. Turk. Dealing with Multi-scale Depth Changes and Motion in Depth Edge Detection. Brazilian Symposium on Computer Graphics and Image Processing (SIBGRAPI 2006), Manaus, Brazil, October 2006. Awarded "One of the Best Image Processing and Computer Vision Papers".
94. H. Guan, J. Chang, L. Chen, R. Feris, and M. Turk. Multi-view Appearance-based 3D Hand Pose Estimation. IEEE Workshop on Vision for Human Computer Interaction (in conjunction with CVPR 2006), New York, NY, June 2006.

95. H. Guan, R. Feris, and M. Turk. The Isometric Self-Organizing Map for Hand Pose Estimation. International Conference on Face and Gesture Recognition, Southampton, UK, 2006.
96. R. Feris, R. Raskar, L. Chen, K. Tan, and M. Turk. Discontinuity Preserving Stereo with Small Baseline Multi-Flash Illumination. International Conference on Computer Vision (**ICCV 2005, Oral Presentation**), Beijing, China, 2005.
97. Y. Zana, R. Cesar, R. Feris, and M. Turk. Face Verification in Polar Frequency Domain: a Biologically Motivated Approach. International Symposium on Visual Computing, Lake Tahoe, NV, 2005.
98. R. Feris, R. Raskar, K. Tan, and M. Turk. Specular Reflection Reduction with Multi-Flash Imaging. Brazilian Symposium on Computer Graphics and Image Processing (SIBGRAPI 2004), Curitiba, Brazil, October 2004. Also accepted as a poster in SIGGRAPH 2004.
99. R. Feris, M. Turk, R. Raskar, K. Tan, and G. Ohashi. Exploiting Depth Discontinuities for Vision-based Fingerspelling Recognition. IEEE Workshop on Real-Time Vision for Human-Computer Interaction (in conjunction with CVPR 2004), Washington DC, USA, June 2004.
100. K. Tan, J. Kobler, R. Feris, P. Dietz, and R. Raskar. Shape Enhanced Surgical Visualizations and Medical Illustrations with Multi-flash Imaging. International Conference on Medical Imaging Computing and Computer Assisted Intervention (MICCAI 2004), Rennes, France 2004.
101. C. Hu, Y. Chang, R. Feris, and M. Turk. Manifold Based Analysis of Facial Expression. IEEE Workshop on Face Processing in Video (in conjunction with CVPR 2004), Washington DC, USA, June 2004.
102. C. Hu, R. Feris, and M. Turk. Real-time View-Based Face Alignment Using Active Wavelet Networks. Workshop on Analysis and Modeling of Faces and Gestures (in conjunction with ICCV 2003), Nice, France, October 2003.
103. C. Hu, R. Feris, and M. Turk. Active Wavelet Networks for Face Alignment. British Machine Vision Conference (BMVC 2003), Norwich, 2003.
104. M. Turk, C. Hu, R. Feris, F. Lashkari, and A. Beall. TLA Based Face Tracking. International Conference on Vision Interfaces, Calgary, Canada, 2002.
105. R. Feris, J. Gemmell, K. Toyama, and V. Krueger. Hierarchical Wavelet Networks for Facial Feature Localization. International Conference on Automatic Face and Gesture Recognition, Washington D.C., USA, May 20-21, 2002.
106. R. Feris, V. Krueger, and R. M. Cesar Jr. Efficient Real-Time Face Tracking in Wavelet Subspace. Workshop on Recognition, Analysis and Tracking of Faces and Gestures in Real-Time Systems (in conjunction with ICCV 2001), Vancouver, BC, 2001.
107. R. Feris and R. M. Cesar Jr. Locating and Tracking Facial Landmarks Using Gabor Wavelet Networks. Lecture Notes in Computer Science, vol. 2013, pp. 311-320. International Conference on Advances in Pattern Recognition, Rio de Janeiro, Brazil, May 2001.
108. V. Krueger and R. Feris. Wavelet Subspace Method for Real-time Face Tracking. Proc. Pattern Recognition. DAGM Symposium, Munich, Germany 2001.

- 109.R. Feris, T. E. Campos, and R. M. Cesar Jr. A Project for Face Recognition from Video Sequences Using GWN and Eigenfeature Selection. Workshop in Artificial Intelligence and Computer Vision, Atibaia, Brazil, November 2000.
- 110.R. Feris and R. M. Cesar Jr. Tracking Facial Features Using Gabor Wavelet Networks. Symposium on Computer Graphics and Image Processing (SIBGRAPI 2000), Gramado, Brazil, October 2000.
- 111.T. Campos, R. Feris, and R. M. Cesar Jr. Improved Face versus Non-Face Discrimination Using Fourier Descriptors through Feature Selection. Brazilian Symposium on Computer Graphics and Image Processing (SIBGRAPI 2000), Gramado, Brazil, October 2000.
- 112.R. Feris, T. Campos, and R. M. Cesar Jr. Detection and Tracking of Facial Features in Video Sequences. Lecture Notes on Artificial Intelligence, vol. 1793, pp. 127-135. Proceedings of MICAI-2000, Acapulco, Mexico, April 2000.
- 113.T. Campos, R. Feris, and R. M. Cesar Jr. Eigenfaces versus Eigeneyes: First Steps Towards Performance Assessment of Representations for Face Recognition. Lecture Notes on Artificial Intelligence, vol. 1793, pp. 193-201. Proceedings of MICAI-2000, Acapulco, Mexico, April 2000.
- 114.R. Feris and W. Lages. Stereo Image Matching Using Correlation and Relaxation Labeling. IV Congress for Scientific Initiation and Postgraduation at Aeronautics Institute of Technology, pp. 193-199, Sao Jose dos Campos-SP, Brazil, October 1998 (in portuguese).

Journal Papers

1. X. Peng , R. Feris, X. Wang, and D. Metaxas. Red-net: A recurrent encoder-decoder network for video-based face alignment. International Journal of Computer Vision (**IJCV**), 2018.
2. M. Merler, D. Joshi, Q. Nguyen, S. Hammer, J. Kent, J. Xiong, M. Do, J. Smith, and R. Feris. Automatic Curation of Sports Highlights using Multimodal Excitement Features. IEEE Transactions on MultiMedia (TMM) 2018.
3. J. Xie, R. Feris, and M.T. Sun. Edge-Guided Single Depth Image Super Resolution. IEEE Transactions on Image Processing (TIP), vol. 25, no. 1, pp. 428-438, 2016.
4. J. Xie, Y. Hsu, R. Feris, and M.T. Sun. Fine registration of 3D point clouds fusing structural and photometric information using an RGB-D camera. Journal of Visual Communication and Image Representation (JVCI), vol. 32, pp. 194-204, 2015.
5. J. Xie, R. Feris, S. Yu, and M.T. Sun. Joint Super Resolution and De-noising from a Single Depth Image. IEEE Transactions on Multimedia (TMM), vol.17, no.9, pp.1525-1537, 2015.
6. R. Feris, B. Siddiquie, J. Petterson, Y. Zhai, A. Datta, L. Brown, and S. Pankanti. Large-Scale Vehicle Detection, Indexing, and Search in Urban Surveillance Videos. IEEE Transactions on Multimedia, 2012.
7. S. Pankanti, L. Brown, J. Connell, A. Datta, Q. Fan, R. Feris, N. Haas, Y. Li, N. Ratha, and H. Thin. Practical Computer Vision: Example Techniques and Challenges. IBM Journal of Research and Development, 201.

8. Y. Tian, R. Feris, H. Liu, A. Hampapur, and M. Sun. Robust Detection of Abandoned and Removed Objects in Complex Surveillance Videos. *IEEE Transactions on Systems, Man, and Cybernetics--Part C: Applications and Reviews*, 2010.
9. H. Liu, R. Feris, V. Krueger, and M.T. Sun. Unsupervised Action Classification Using Space-Time Link Analysis. *Eurasip Journal on Advances in Signal Processing*, 2010.
10. A. Hampapur et al. Analytics Driven Asset Management. *IBM Smart Cities Journal*, 2010.
11. R. Feris, R. Raskar, L. Chen, K. Tan, and M. Turk. Multi-Flash Stereopsis: Depth Edge Preserving Stereo with Small Baseline Illumination. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 30, no. 1, pp. 147-159, 2008.
12. Y. Zana, R. Cesar, R. Feris and M. Turk. Local Approach for Face Verification in Polar Frequency Domain. *Image and Vision Computing Journal*, vol. 24, no. 8, pp. 904-913, 2006.
13. R. Feris, M. Turk, R. Raskar, and K. Tan. Specular Highlights Detection and Reduction with Multi-Flash Photography. *International Journal of the Brazilian Computer Society*, vol. 1, no. 12, pp. 35-42, 2006.
14. Y. Chang, C. Hu, R. Feris, and M. Turk. Manifold-Based Analysis of Facial Expressions. *Image and Vision Computing Journal*, vol.24, no. 6, pp. 605-614, 2006.
15. R. Raskar, K. Tan, R. Feris, J. Kobler, J. Yu, and M. Turk. Harnessing Real-World Depth Edges with Multi-Flash Imaging. *IEEE Computer Graphics and Applications (IEEE CG&A)*, vol. 25, no. 1, pp. 32-38, January 2005.
16. R. Feris, V. Krueger, and R. Cesar. A Wavelet Subspace Method for Real-time Face Tracking. *Journal of Real-time Imaging*, vol. 10, pp. 339-350, 2004.
17. R. Raskar, K. Tan, R. Feris, J. Yu, and M. Turk. Non-photorealistic Camera: Depth Edge Detection and Stylized Rendering using Multi-Flash Imaging. *ACM Transactions on Graphics (SIGGRAPH 2004)*, Vol. 23, Issue 3, August 2004. Also accepted in *SIGGRAPH Emergent Technologies*, 2004.

Book Chapters

1. R. Feris, C. Lampert and D. Parikh. Introduction to Visual Attributes. In *Visual Attributes* (eds: R. S. Feris, C. Lampert, and D. Parikh), Springer, 2017.
2. H. Liu, R. Feris, and M.T. Sun. Benchmarking Datasets for Human Activity Recognition. *Visual Analysis of Humans – Looking at People*, Springer 2011.
3. Y. Zhai, R. Feris, A. Hampapur, S. Russo, and S. Pankanti. Parsing Object Events in Heavy Urban Traffic. *Object Tracking*, Intech, 2011.
4. H. Liu, M.T. Sun, and R. Feris. Video Activity Recognition. *Multimedia Analysis, Processing and Communication*, Z. Li, J. Kacprzyk, D. Tao, E. Izquierdo, W. Lin, and H. Wang ed., Springer, 2010.
5. R. Feris, A. Hampapur, Y. Zhai, R. Bobbitt, L. Brown, D. Vaquero, Y-L. Tian, H. Liu, and M-T. Sun. Case Study: IBM Smart Surveillance System. *Intelligent Video Surveillance: Systems and Technology*, by Taylor & Francis Group, LLC, 2009.

6. Y. Zhai, R. Feris, L. Brown, R. Bobbitt, A. Hampapur, S. Pankanti, Q. Fan, A. Yanagawa, Y. Tian, and S. Velipasalar. Composite Event Detection in Multi-Camera and Multi-Sensor Surveillance Networks. *Multi-Camera Networks: Concepts and Applications*, by Elsevier, 2009.
7. Y. Tian, R. Feris, L. Brown, D. Vaquero, Y. Zhai, and A. Hampapur. Multi-Scale People Detection and Motion Analysis for Video Surveillance. *Machine Learning for Human Motion Analysis: Theory and Practice*, IGI Global, 2009.
8. D. Vaquero, R. Feris, L. Brown, A. Hampapur, and Matthew Turk. Attribute-based People Search. *Intelligent Video Surveillance: Systems and Technology*, by Taylor & Francis Group, LLC, 2009.
9. Y. Tian, A. Hampapur, L. Brown, R. Feris, M. Lu, A. Senior, C. Shu, and Y. Zhai. Event Detection, Query, and Retrieval for Video Surveillance. *Artificial Intelligence for Maximizing Content Based Image Retrieval*, 2008.
10. R. Feris, M. Turk, R. Raskar, K. Tan, and G. Ohashi. Recognition of Isolated Fingerspelling Gestures Using Depth Edges. B. Kisacanin, V. Pavlovic and T. Huang (eds.), *Real-time Vision for Human-Computer Interaction*, Springer-Verlag, 2005.

Theses and Reports

1. M. Merler, N. Ratha, R. Feris, and J. Smith. Diversity in Faces. arXiv 2019.
2. B. Cheng, Y. Wei, R. Feris, J. Xiong, W. Hwu, T. Huang, and H. Shi. Decoupled Classification Refinement: Hard False Positive Suppression for Object Detection. arXiv 2018.
3. R. Feris. Detection and Analysis of Depth Discontinuities with Lighting and Viewpoint Variation. PhD thesis, University of California, Santa Barbara, 2006.
4. R. Feris, J. Gemmell, K. Toyama, and V. Krueger. Facial Feature Detection using a Hierarchical Wavelet Face Database. Microsoft Research Technical Report MSR-TR-2002-05, January 2002.
5. R. Feris. Efficient Real-time Face Tracking in Wavelet Subspace. MSc. thesis, University of Sao Paulo, May 2001. Awarded as the Second Best Computer Science MSc. Thesis in Brazil (Nationwide Competition, 2002).

Demos and Posters

1. D. Joshi, M. Merler, Q. Nguyen, S. Hammer, J. Kent, J. Smith, and R. Feris. IBM High-Five: Highlights from Intelligent Video Engine. *ACM Multimedia*, 2017.
2. F. Yu, L. Cao, R. Feris, J. Smith, and S. F. Chang. Designing Category-level Attributes for Discriminative Visual Recognition. *Greater New York Area Workshop on Multimedia and Vision*, 2013. (Best Poster Award)
3. IBM Centennial Demo – Traffic Analysis Using Video Analytics. Lincoln Center, Manhattan, New York 2011.

4. R. Feris, L. Brown, Q. Fan, A. Datta, S. Pankanti, A. Hampapur, Y. Zhai, R. Bobbitt, R. Kjeldsen, M. Lu, C. Shu, S. Russo, and C. Milite. IBM Smart Surveillance System: Searchable Video Analytics for Urban Environments and Retail Stores. CVPR DEMO, 2010.
5. A. Agrawal, V. Branzoi, R. Chellappa, R. Feris, R. Raskar, K. Tan, and M. Turk. Depth Edges in Real-Time Using a Multi-Flash Camera. CVPR DEMO, 2005.
6. R. Raskar, K. Tan, R. Feris, J. Yu, and M. Turk. Non-photorealistic Camera: Automatic Stylization with Multi-Flash Imaging. SIGGRAPH Emergent Technologies, 2004.
7. R. Feris, R. Raskar, K. Tan, and M. Turk. Specular Reflection Reduction Using a Multi-Flash Camera. SIGGRAPH poster, 2004.

Issued Patents

1. 9,495,599 - Determination of train presence and motion state in railway environments
2. 9,477,890 - Object detection using limited learned attribute ranges
3. 9,471,852 - User-configurable settings for content obfuscation
4. 9,460,361 - Foreground analysis based on tracking information
5. 9,460,349 - Background understanding in video data
6. 9,443,148 - Visual monitoring of queues using auxiliary devices
7. 9,430,874 - Estimation of object properties in 3D world
8. 9,424,659 - Real time processing of video frames
9. 9,396,548 - Multi-cue object detection and analysis
10. 9,342,594 - Indexing and searching according to attributes of a person
11. 9,330,312 - Multispectral detection of personal attributes for video surveillance
12. 9,330,111 - Hierarchical ranking of facial attributes
13. 9,322,647 - Determining camera height using distr of object heights and object image heights
14. 9,299,162 - Multi-mode video event indexing
15. 9,280,833 - Topology determination for non-overlapping camera network
16. 9,262,445 - Image ranking based on attribute correlation
17. 9,251,425 - Object retrieval in video data using complementary detectors
18. 9,245,186 - Semantic parsing of objects in video
19. 9,224,049 - Detection of static object on thoroughfare crossings
20. 9,224,046 - Multi-view object detection using appearance model transfer from similar scenes
21. 9,165,375 - Automatically determining field of view overlap among multiple cameras
22. 9,134,399 - Attribute-based person tracking across multiple cameras
23. 9,104,919 - Multi-cue object association
24. 9,082,201 - Surface contamination determination system
25. 9,069,104 - Pathway management using model analysis and forecasting
26. 9,058,669 - Incorporating video meta-data in 3D models
27. 8,948,454 - Boosting object detection performance in videos
28. 8,934,670 - Real time processing of video frames for triggering an alert
29. 8,837,776 - Rule-based combination of a hierarchy of classifiers for occlusion detection
30. 8,824,791 - Color correction for static cameras
31. 8,811,663 - Object detection in crowded scenes
32. 8,774,532 - Calibration of video object classification
33. 8,675,917 - Abandoned object recognition using pedestrian detection
34. 8,620,026 - Video-based detection of multiple object types under varying poses
35. 8,488,881 - Object segmentation at a self-checkout
36. 8,483,481 - Foreground analysis based on tracking information
37. 8,249,301 - Video object classification
38. 8,170,276 - Object detection system based on a pool of adaptive features

39. 8,107,678 - Detection of abandoned and removed objects in a video stream
40. 7,738,725 - Stylized rendering using a multi-flash camera

Invited Talks

- Keynote speaker, ICML LatinX in AI workshop, 2020. “Dynamic Neural Networks for Efficient Image and Video Classification”
- Keynote speaker, CVPR Workshop on Diagram Image Retrieval and Analysis, 2020. “Visual Learning Beyond Natural Images”
- Invited speaker, What’s Next in AI – AI We can Scale, MIT-IBM Watson AI Lab event. “Learning to Learn”
- Invited speaker, NeurIPS Workshop on Energy Efficient Machine Learning and Cognitive Computing, 2019. “Adaptive (Dynamic) Neural Networks for Efficient Inference”
- Invited speaker, CVPR Workshop on Learning from Imperfect Data, 2019. “Learning More from Less: Weak Supervision and Beyond”
- Keynote speaker, CVPR Workshop on Understanding Subjective Attributes of Data, 2019. “Is It All Relative? Interactive Fashion Search based on Relative Natural Language Feedback”
- Invited speaker, CVPR Workshop on Energy Efficient Machine Learning and Cognitive Computing, 2019. “Speeding up Deep Neural Networks with Adaptive Computation and Efficient Multi-Scale Architectures”
- Keynote speaker, Conference on Graphics, Patterns and Images (SIBGRAPI 2016), Sao Jose dos Campos, Brazil, 2016. “Representation Learning Beyond Human Labels: Practical Applications in Visual Analysis of People”
- Keynote speaker, International Joint Conference on Artificial Intelligence (IJCAI) BeyondLabeler Workshop, New York City, USA, 2016. “Representation Learning Beyond Human Labels: Practical Applications in Visual Analysis of People”
- Invited Speaker, NYC Computer Vision Meetup, 2015. “Deep Domain Adaptation for Describing People Based on Fine-grained Attributes”
- Invited Speaker, Workshop on Recent Trends in Computer Vision, Boston, USA, 2015. “Deep Domain Adaptation for Describing People Based on Fine-grained Attributes”.
- Invited Speaker, International Conference on Multimedia Retrieval (ICMR) Practitioner’s day, Glasgow, Scotland, 2014. “Multimedia Analytics at IBM Research”
- NYU Center for Urban Science and Progress, New York City, 2014. “Smart Surveillance: Towards Visual Search based on Fine-Grained Semantic Attributes”. Host: Claudio Silva
- Invited Speaker, CVPR Workshop on Fine-Grained Visual Categorization, Portland, Oregon, 2013. “Fine-Grained Attribute Classification and Search: Two Case Studies and Lessons Learnt”
- IBM Research, Brazil, 2013. “Learning Visual Semantics”. Host: Andrea Mattos

- Invited Speaker, National Academy of Engineering JAF0E meeting, Irvine, California, 2012. “Intelligent Video Surveillance”
- University of Delaware, 2012. “Attribute-based Object Retrieval”. Host: Jingyi Yu
- Invited Speaker, ICMR 2011 Practitioner’s day, Trento, Italy 2011. “Attribute-based Object Retrieval”
- City College of New York, May 2010. “Searching for People and Vehicles in Urban Environments”. Host: Yingli Tian
- Cognitec, Dresden, Germany, September 2008. “The IBM Smart Surveillance System”. Host: Roger Kelesoglu
- Invited speaker, IBM Conference on Maturing and Leveraging Biometrics Analytics, 2008. “Person Identification through Search Based on Physical Attributes”
- University of Washington, 2007. “The IBM Smart Surveillance System”. Host: Linda Shapiro
- Princeton University, 2007. “Computer Vision for Digital Photography and Smart Surveillance”. Host: Fei-Fei Li
- Kodak Research, Intelligent Systems Lab, Rochester, NY, 2007. “Multi-Flash Photography and Applications”. Host: Jiebo Luo
- HP Laboratories, Computer Vision and Graphics group, Palo Alto, CA, 2007. “Detection and Modeling of Depth Discontinuities with Lighting and Viewpoint Variation”. Host: Bruce Culbertson
- Adobe Advanced Technology Labs, Seattle, 2006. “Discontinuity Detection and Modeling with Lighting and Viewpoint Variation” (also presented at University of Washington). Host: David Salesin
- Jet Propulsion Laboratory, Pasadena, CA, 2006. “Detection and Analysis of Depth Discontinuities in Computer Vision”. Host: Larry Matthies
- IBM T.J. Watson Research, Hawthorne, NY, 2005. “What Can We Do with a Multi-Flash Camera?” Host: Deepak Turaga (as part of Emerging Leaders in Multimedia Seminar)
- Microsoft Research, Redmond, 2001. “A Wavelet Subspace Method for Real-time Face Tracking”. Host: Kentaro Toyama

Academic Experience

- **Adjunct Associate Professor, Columbia University**
 - Instructor: Visual Recognition and Search, Spring 2014 (Graduate course)
<http://rogerioferis.com/VisualRecognitionAndSearch2014/home.html>
 Evaluation: 5.00/5.00
 - Instructor: Visual Recognition and Search, Spring 2013 (Graduate course)
<http://rogerioferis.com/VisualRecognitionAndSearch2013/home.html>
 Evaluation: 4.56/5.00

- Instructor: Automatic Video Surveillance, Spring 2008 (Graduate Course)
<http://www.andrewsenior.com/technical/surveillanceclass/>
- **Affiliate Associate Professor, University of Washington**
 - Co-advised the PhD dissertations of Haowei Liu and Jun Xie (both working at Google now) and was involved in several other academic activities (such as selection of PhD students)
- **Short Courses / Tutorials**
 - ICCV 2019 Tutorial on Visual Learning with Limited Labeled Data
<https://sites.google.com/view/learning-with-limited-data/home>
with Kate Saenko, Leonid Karlinsky, Ross Girshick, Judy Hoffman, Kevin Swersky, and Lorenzo Torresani
 - ICCV 2015 Tutorial on Tools for Efficient Object Detection
<http://mp7.watson.ibm.com/ICCV2015/ObjectDetectionICCV2015.html>
with Piotr Dollar, Xiaoyu Wang, Kaiming He, Ross Girshick, Rodrigo Benenson, and Jan Hosang
 - CVPR 2014 Tutorial on Learning Visual Semantics: Models, Massive Computation, and Innovative Applications
<http://mp7.watson.ibm.com/LearningVisualSemantics/index.htm>
with Liangliang Cao, John Smith, and Shih-Fu Chang
 - CVPR 2012 Tutorial on Looking at People: Benchmarking Human Activity Recognition
with Haowei Liu and Ming-Ting Sun

Professional Activities

- **Program Chair**
 - IEEE Winter Conference on Applications of Computer Vision (WACV) 2017
 - 1st GNY Area Vision and Multimedia Meeting, 2011.
- **Area Chair**
 - Conference on Computer Vision and Pattern Recognition (CVPR) - 2015, 2016, 2017, 2020, 2021
 - Neural Information Processing Systems (NeurIPS) - 2018, 2019, 2020
 - International Conference on Learning Representations (ICLR) - 2021
 - European Conference on Computer Vision (ECCV) - 2020
 - International Conference on Machine Learning (ICML) - 2021
 - International Conference on Computer Vision (ICCV) - 2015
 - ACM Multimedia - 2017
 - International Symposium on Visual Computing (ISVC) - 2015
- **Organizing Committee**
 - CVPR Workshop on Dynamic Neural Networks Meets Computer Vision, 2021
 - CVPR Workshop on Neural Architecture Search, 2021
 - CVPR Workshop on Visual Learning with Limited Labels: Zero-Shot, Few-Shot, Any-Shot, and Cross-Domain Few-Shot Learning, 2020

- CVPR Workshop on Computer Vision for Fashion, Art, and Design, 2020
 - CVPR Workshop on Neural Architecture Search and Beyond for Representation Learning, 2020
 - CVPR Workshop on Diagram Image Retrieval and Analysis, 2020
 - CVPR Workshop on Fair, Data-Efficient and Trusted Computer Vision, 2020
 - CVPR Workshop on Skin Image Analysis, 2020 (Steering Committee)
 - ECCV Workshop on Multimodal Video Analysis, 2020
 - CVPR Workshop on Bias Estimation in Face Analytics, 2019
 - CVPR Workshop on Skin Image Analysis, 2019 (Steering Committee)
 - ICCV Workshop on Linguistics Meets Image and Video Retrieval, 2019
 - ICCV Workshop on Moving Cameras: from Body Cameras to Drones, 2019
 - ICCV Workshop on Computer Vision for Fashion, Art, and Design, 2019 (Steering Committee)
 - ICCV Workshop on Multimodal Video Analysis and Moments in Time Challenge, 2019
 - ICML Workshop on Towards learning with limited labels: Equivariance, Invariance, and Beyond, 2018
 - ECCV Workshop on Bias Estimation in Face Analytics, 2018
 - ECCV Workshop on Computer Vision for Fashion, Art, and Design, 2018
 - CVPR Workshop on Moving Cameras: from Body Cameras to Drones, 2016
 - ECCV Workshop on Parts and Attributes, 2014
 - ECCV Workshop on Parts and Attributes, 2012
 - GNY Area Vision and Multimedia Meeting, 2012
 - ECCV Workshop on Parts and Attributes, 2010
 - IBM Watson Emerging Leaders in Multimedia Workshop, 2006 and 2007
 - I Workshop on Artificial Intelligence and Computer Vision, 1997
- **Industry chair**
 - Corporate Relations Chair, CVPR 2018
 - Industry and Practitioner Chair, ICMR 2016
 - Corporate Relations Chair, FG 2011
- **LDV Vision Summit:** judge for the Entrepreneurial Computer Vision Challenges (<http://www.ldv.co/visionsummit/2015/speakers>), 2015
- **Award Committee Member**
 - City University of Hong Kong (CUHK) International Selection Committee for PhD thesis Award, 2015
 - Member of the committee for selection of the best paper award – International Conference in Image Processing (ICIP 2008)
- Watson Chair of the Multimedia Professional Interest Community (2013 – 2017)
 - Doctoral Consortium Co-Chair, IEEE Conference on Automatic Face and Gesture Recognition (FG 2011), Santa Barbara, USA, 2011
 - Grant reviewer: Expert proposal evaluator for FIT-IT, Vienna, Austria, 2011.
 - Department Intern Coordinator, IBM Research, 2010
 - Seminar PIC Coordinator, IBM Research, 2010-2011
 - Program Committee Member: CVPR (2007-2014), ICCV (2007, 2009, 2011), ECCV (2008,2010,2012), AVSS (2008), WACV (2011), ICIP (2009), ICPR (2010,2014), ACCV (2007,2009), Visual Surveillance Workshop (2009,2010) and others

- Reviewer for major conferences and journals in computer vision/graphics, including ACM SIGGRAPH, IEEE Transactions on PAMI, Pattern Recognition Letters, IVC Journal, CVIU Journal, CVPR, ICCV, Face and Gesture Recognition, Eurographics, ISMAR, etc.

Media Press Coverage

- Shrinking deep learning's carbon footprint. MIT News, 2020.
- Toward a machine learning model that can reason about everyday actions. MIT News, 2020.
- IBM researchers develop a pair of low-power, high-performance computer vision systems. VentureBeat, 2018.
- Coffee delivery drone patented by IBM. BBC News, 2018.
- Enjoy Those U.S. Open Highlights. A Computer Picked Them for You. New York Times, 2017.
- How an IBM Computer Picks U.S. Open Highlights. Fortune, 2017.
- IBM's Watson Serves Up This Year's U.S. Open Highlights. NBC News, 2017.
- How Artificial Intelligence May be Tennis's Ace for Serving Grand Slam Fans. Newsweek, 2017.
- IBM's Watson is creating US Open tennis highlight videos. Engadget, 2017.
- IBM uses AI to serve up Wimbledon highlights. CNet, 2017.
- IBM To Provide Wimbledon Highlights Using Artificial Intelligence. SportTechie, 2017.
- IBM Watson is creating highlight reels at the Masters. ZDNet, 2017.
- Fighting terrorism in New York City. CBS 60 Minutes, 2011. (starting at minute 7:20)
- ABC7 puts video analytics to the test. ABC News, 2010.
- Cameras help confirm Scott suicide ruling. ABC News, 2009.
- Intelligent Iris. ABC News, 2008.
- The Nonphotorealistic camera. SlashDot, 2004.
- Master's dissertation / Brazilian Media: My work on face tracking was featured at Jornal do SBT and TV USP.
- High-School Project / Brazilian Media: In 1993, during my technical high-school, I developed a software for constructing facial sketches from witnesses' descriptions (coding in assembly for accessing the video memory card, dealing with bit planes, etc.) My work was reported in several newspapers, magazines, and television programs, including Ciencia Viva and RBS TV.

Ph.D. Thesis Committees

- Karim Ahmed, Dartmouth College, 2019
- Brendan Jou, Columbia University, 2016
- Xiaodong Yang, City University of New York, 2015
- Shizhi Chen, City University of New York, 2013
- Eanes Pereira, Universidade Federal de Campina Grande, 2012
- Haowei Liu, University of Washington, 2011