

Sangmin Oh, Ph.D.

(E-mail / Homepage / Mobile) sangmin.oh1@gmail.com / sangminoh.org / 404-547-8622

Summary

Equally comfortable in conducting R&D, and program management. R&D interest includes designing, building, and evaluating real-world intelligent perceptual systems based on sound theoretical background under complex real-world constraints where data is very large, noisy, partially missing, and highly uncertain. Areas of expertise include computer vision, machine learning, multimedia, embedded systems, social network analysis, wearable computing, time-series data modeling, and robotics. Experienced in solving challenging problems across both fundamental and applied research, technology transfer, fast-paced development, leading R&D teams, and designing and testing end-to-end systems.

Education

- Ph.D in Computer Science, Georgia Institute of Technology, May. 2009.
- M.S. in Computer Science, Georgia Institute of Technology, Jan 2009.
- B.S. in Computer Science with *cum laude*, Seoul National University, Korea, Feb. 2003.

Employment History

- Employer: Nvidia, Senior Computer Vision Software Engineer, Dec 2014 - Present
Job Description: Computer vision software R&D for embedded and mobile systems
- Employer: Kitware, Research Engineer, Sep 2009 - Dec 2014
Job Description: R&D in computer vision, robotic perception, and social network analysis, sponsored by US Government and commercial companies. Published +20 research articles and led collaboration with top-tier academic institutions and industrial research labs. Mentored 4 PhD and 2 MS level interns.
- Employer: Georgia Institute of Technology, Sep 2003 - Sep 2009
Advisers: Frank Dellaert and James M. Rehg
Job Description: Development and evaluation of state-of-the-art algorithms for the problems in the domain of computer vision, robotics, and wearable computing.
- Employer: Microsoft Research, Redmond, WA, Research Intern, Summer 2007 & Summer 2006
Collaborators/Mentors: Paul Viola, Cha Zhang, and Nebojsa Jojic.
Job Description: Worked as research interns for research, development and evaluation of algorithms for rapid object detection and robotic mapping/navigation.
- Employer: Republic of Korea Navy, Sep 1997 - Dec 1999.
Job Description: Served military service for 2.4 years.

Project Experience

For representative project and research topics, please see below:

- **Recognition of Objects, Scenes, and Events from Images and Videos**
 - Conducted and led research and development on large-scale visual recognition from images and videos collected from web, social network, and surveillance cameras (both ground and aerial). Developed and owned algorithms softwares for visual classifiers, machine learning, feature fusion, feature encoding, and content summarization. Played key roles in participating TRECVID multimedia event detection and recounting for 3 years (2011-2013), and creating large-scale datasets for surveillance (CVPR 11).
 - Led research collaboration with 9 world-class universities.
 - Managed funding total of approximately 6M USD, funded by DARPA and IARPA.
 - 20 Publications and 2 Patents.

- **Robotic Perception, Geometry, Localization, Control**
 - Conducted research on geometric scene understanding and object recognition to improve robot navigation and robotic interaction with humans.
 - Led research collaboration with 3 world-class universities.
 - Manged funding total of approximately 2M USD, funded by DARPA.
 - 6 Publications
- **Time-series Modeling for Activity Recognition from Videos and Wearable Motion Sensors**
 - Conducted research on developing probabilistic temporal models for visual and motion data streams for activity recognition for animals and humans.
 - 4 Publications

Publications

1. **Compositional Models for Video Event Detection: A Multiple Kernel Learning Latent Variable Approach**, Arash Vahdat, Kevin Cannons, Greg Mori, Sangmin Oh and Ilseo Kim, submitted to IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), under review.
2. **Towards visual analysis of unconstrained images in social forums: Studies on concept detection and personalized economy of images in social networks**, Sangmin Oh, Eric Smith, Yiliang Xu, and Anthony Hoogs, to appear in IEEE International Workshop on Applied Imagery and Pattern Recognition, 2014.
3. **Personalized Economy of Images in Social Forums: An Analysis on Supply, Consumption, and Saliency**, Sangmin Oh, Megha Pandey, Ilseo Kim, Anthony Hoogs, and Jeff Baumes, in Proceedings of International Conference on Pattern Recognition (ICPR), 2014, *winner of Best Industry-related Paper Award*.
4. **Multimedia Event Detection with Multimodal Feature Fusion and Temporal Concept Localization**, Sangmin Oh, Scott McCloskey, Ilseo Kim, Arash Vahdat, Kevin Cannons, Hossein Hajimirsadeghi, Greg Mori, Amitha A.G. Perera, Megha Pandey, and Jason Corso. in Machine Vision and Applications (MVA) Special issue on Multimedia Event Detection, 25:49-69, Springer, 2014.
5. **Compositional Models for Video Event Detection: A Multiple Kernel Learning Latent Variable Approach**, Arash Vahdat, Kevin Cannons, Greg Mori, Sangmin Oh and Ilseo Kim, in Proceedings of International Conference on Computer Vision (ICCV) 2013.
6. **Learning Non-linear Calibration for Score Fusion with Applications to Image and Video classification**, Tianyang Ma, Sangmin Oh, A.G. Amitha Perera, and Longin Jan Latecki, in 2nd International Workshop on Large-Scale Video Search and Mining, in conjunction with International Conference on Computer Vision (ICCV) 2013.
7. **Segmental Multi-way Local Pooling for Video Recognition**, Ilseo Kim, Sangmin Oh, Arash Vahdat, Kevin Cannons, A.G. Amitha Perera, and Greg Mori, in Proceedings of ACM Multimedia 2013.
8. **TRECVID 2013 GENIE: Multimedia Event Detection and Recounting**, Sangmin Oh, A.G. Amitha Perera, Ilseo Kim, Megha Pandey, Kevin Cannons, Hossein Hajimirsadeghi, Arash Vahdat, Greg Mori, Ben Miller, Scott McCloskey, You-Chi Cheng, Zhen Huang, Chin-Hui Lee, Chenliang Xu, Rohit Kumar, Wei Chen, Jason Corso, L. Fei-Fei, Daphne Koller, Vignesh Ramanathan, Kevin Tang, Armand Joulin, and Alexandre Alahi, in NIST TrecVID workshop 2013.
9. **System and Algorithms on Detection of Objects Embedded in Perspective Geometry using Monocular Cameras**, Yiliang Xu, Sangmin Oh, Fan Yang, Zhuolin Jiang, Anthony Hoogs and Larry Davis, in IEEE International Conference on Advanced Video and Signal-Based Surveillance (AVSS), 2013.
10. **A Minimum Error Vanishing Point Detection Approach for Uncalibrated Monocular Images of Man-made Environments**, Yiliang Xu, Sangmin Oh, Anthony Hoogs, in IEEE International Conference on Computer Vision and Pattern Recognition (CVPR), 2013.
11. **Automatic Building Exterior Mapping Using Multilayer Feature Graphs**, Yan Lu, Dezhen Song, Yiliang Xu, A.G. Amitha Perera, Sangmin Oh, in 9th IEEE International Conference on Automation Science and Engineering (CASE) 2013.
12. **A Videography Analysis Framework for Video Retrieval and Summarization**, Sangmin Oh, Kang Li, A.G. Amitha Perera, Yun Fu, in Proceedings of IEEE British Machine Vision Conference (BMVC) 2012 (oral).

13. **Explicit Performance Metric Optimization for Fusion-based Video Retrieval**, Ilseo Kim, Sangmin Oh, Byungki Byun, A.G. Amitha Perera, Chin-Hui Lee, in Proceedings of Workshop on Information Fusion in Computer Vision for Concept Recognition in conjunction with European Conference on Computer Vision (ECCV), 2012.
14. **TRECVID 2012 GENIE: Multimedia Event Detection and Recounting**, Perera A.G.A., Oh S., Pandey M., Ma T., Hoogs A., Vahdat A., Cannons K., Hajimirsadeghi H., Mori G., McCloskey S., Miller B., Venkatesh S., Davalos P., Das P., Xu C., Corso J., Srihari R., Kim I., Cheng Y-C., Huang Z., Lee C-H., Tang K., Li F-F., Koller D., in NIST TRECVID Workshop 2012.
15. **Per-exemplar Fusion Learning for Video Retrieval and Recounting**, Ilseo Kim, Sangmin Oh, A.G.Amitha Perera, Chin-Hui Lee, in Proceedings of IEEE 2012 International Conference on Multimedia and Expo (ICME), 2012 (oral).
16. **A Large-scale Benchmark Dataset for Event Recognition in Surveillance Video**, S. Oh, A. Hoogs, A. Perera, N. Cuntoor, C-C. Chen, J.T. Lee, S. Mukherjee, J.K. Aggarwal, H. Lee, L. Davis, E. Swears, X Wang, Qiang Ji, K. Reddy, M. Shah, C.Vondrick, H. Pirsiavash, D. Ramanan, J. Yuen, A. Torralba, Bi Song, A. Fong, A. Roy-Chowdhury, and M. Desai, in Proceedings of IEEE 2011 International Conference on Computer Vision and Pattern Recognition (CVPR), 2011.
17. **GENIE TRECVID 2011 Multimedia Event Detection: Late-Fusion Approaches to Combine Multiple Audio-Visual Features**, A.G.A. Perera, S. Oh, M. Leotta, I. Kim, B. Byun, C-H. Lee, S. McCloskey, J. Liu, B. Miller, Z.F. Huang, A. Vahdat, W. Yang, G. Mori, K. Tang, D. Koller, L. Fei-Fei, K. Li, G. Chen, J. Corso, Y. Fu, R. Srihari, in NIST TRECVID Workshop 2011.
18. **Augmenting aerial earth maps with dynamic information from videos**, K. Kim, S. Oh, J. Lee, and I. Essa, in Virtual Reality Vol 15, Issue 2-3, pages 185-200, 2011, Springer London.
19. **Discriminative Adaptation of Codebooks for Large-Scale Video Retrieval with Varying Content**, Sangmin Oh, Amitha Perera, Anthony Hoogs, presented in Snowbird Learning Workshop, 2011.
20. **Content-based Retrieval of Functional Objects in Video using Scene Context**. Sangmin Oh, Hoogs A., Turek M., Collins R., in Proceedings of European Conference on Computer Vision (ECCV), 2010.
21. **Unsupervised Learning of Activities in Video using Scene Context**. Sangmin Oh, Hoogs A., in Proceedings of 2010 International Conference on Pattern Recognition (ICPR), Aug-2010.
22. **Temporal Causality for the Analysis of Visual Events**. Prabhakar K., Sangmin Oh, Wang P., Abowd G., Rehg J., in Proceedings of IEEE 2010 International Conference on Computer Vision and Pattern Recognition (CVPR), 2010.
23. **Learning Visibility of Landmarks for Vision-based Localization**. Alcantarilla P.F., Sang Min Oh, Mariottini G.L., Bergasa L.M., Dellaert F., in Proceedings of IEEE International Conference on Robotics and Automation (ICRA), 2010.
24. **Content-based Retrieval of Functional Objects in Video using Scene Context**, Sangmin Oh and Anthony Hoogs, presented in Snowbird Learning Workshop, 2010.
25. **Augmenting Earth Maps with Dynamic Information**. Kihwan Kim, Sang Min Oh, Jeongkyu Lee, Irfan Essa, In Proceedings of IEEE and ACM 2009 International Symposium on Mixed and Augmented Reality (ISMAR).
26. **Switching Linear Dynamic Systems with Higher-order Temporal Structure**, PhD Thesis, Sang Min Oh, May 2009.
Committee members : Prof. Frank Dellaert, Prof. James M. Rehg, Prof. Aaron Bobick, Prof. Irfan Essa, and Prof. Padhraic Smyth.
27. **Learning and Inferring Motion Patterns using Parametric Segmental Switching Linear Dynamic Systems**, Sang Min Oh and James M. Rehg and Tucker Balch and Frank Dellaert, in International Journal of Computer Vision Special Issue on Learning for Vision, 2008 (IJCV 2008), May 2008. Vol.77(1-3). Pages 103-124.
28. **Traversability Classification for UGV Navigation : A Comparison of Patch and Superpixel Representation**, Dongshin Kim, Sang Min Oh, James M. Rehg, in Proceedings of 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).
29. **Parameterized duration modeling for switching linear dynamic systems**, Sang Min Oh, James M. Rehg, and Frank Dellaert, in Proceedings of 2006 IEEE International Conference on Computer Vision and Pattern Recognition (CVPR).

30. **Traversability classification using unsupervised on-line visual learning for outdoor robot navigation**, Dongshin Kim, Jie Sun, Sang Min Oh, James M. Rehg, and Aaron F. Bobick, in Proceedings of 2006 IEEE International Conference on Robotics and Automation (ICRA).
31. **Learning and Inferring Motion Patterns using Parametric Segmental Switching Linear Dynamic Systems**, Sang Min Oh, James M. Rehg, Tucker Balch, and Frank Dellaert, Georgia Tech GVU Report GIT-GVU-06-02 (2006).
32. **Automatic Acquisition of 4D Urban Models and Proactive Auditory Service for Enhanced User Experience**, Sang Min Oh, Grant Schindler, Frank Dellaert, Bruce N. Walker, Jefferey Lindsay, in Young Investigators Forum in Culture and Technology, Daejon, Republic of Korea, 2006.
33. **Learning and Inference in Parametric Switching Linear Dynamical Systems**, Sang Min Oh, James M. Rehg, Tucker Balch, and Frank Dellaert, in Proceedings of IEEE International Conference on Computer Vision (ICCV), 2005.
34. **Data-Driven MCMC for Learning and Inference in Switching Linear Dynamic Systems**, Sang Min Oh, James M. Rehg, Tucker Balch, and Frank Dellaert, Proceedings of Twentieth National Conference on Artificial Intelligence (AAAI), 2005.
35. **Segmental Switching Linear Dynamic Systems**, Sang Min Oh, James M. Rehg, and Frank Dellaert, Georgia Tech CoC Report GIT-CC-05-13 (2005).
36. **A Variational inference method for Switching Linear Dynamic Systems**, Sang Min Oh, Ananth Ranganathan, James M. Rehg, and Frank Dellaert, Georgia Tech GVU Report GIT-GVU-05-16 (2005).
37. **Mixture Trees for Modeling and Fast Conditional Sampling with Applications in Vision and Graphics**, Frank Dellaert, Vivek Kwatra, and Sang Min Oh, in Proceedings of IEEE International Conference on Computer Vision and Pattern Recognition.
38. **Map-Based Priors for Localization**, Sang Min Oh, Sarah Tariq, Bruce Walker, and Frank Dellaert, in Proceedings of IEEE/RSJ International Conference on Intelligent Robotics and Systems.

Tutorials

- Tutorial “Emerging Topics in Human Activity Recognition” with Michael Ryoo, Ivan Laptev and Greg Mori, at International Conference on Computer Vision and Pattern Recognition, Columbus, Ohio, 2014.
- Tutorial “Activity Recognition for Visual Surveillance” with Michael Ryoo, Anthony Hoogs, and Arslan Basharat, at IEEE International Conference on Advanced Video and Signal-Based Surveillance, Beijing, China, 2012.

Awards

- Best Industry-related Paper Award at International Conference on Pattern Recognition (ICPR), 2014.
- Samsung Lee Kun Hee Scholarship Fellow, total 400K USD, Fall 2003 - Spring 2007.
- Highest Honors with full scholarship, Seoul National University, Korea, 2001-2003
- Research Science Institute (RSI) fellowship awardee, 1994.

Synergistic Activities

- Reviewers for the following journals
 - IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI),
 - IEEE Transactions on Image Processing (T-IP)
 - IEEE Transactions on Signal Processing (T-SP)
 - Computer Vision and Image Understanding (CVIU), Elsevier.
 - Image and Vision Computing (IVC)

- Program Committee member, reviewers for the following conferences:
 - International Conference on Computer Vision and Pattern Recognition (CVPR), 2005-2014.
 - International Conference on Computer Vision (ICCV), 2005-2013.
 - European Conference on Computer Vision (ECCV), 2006-2014.
 - AAAI Conference on Artificial Intelligence (AAAI), 2005, 2006, 2007, 2012.
 - International Conference on Robotics and Automation (ICRA), 2005, 2006, 2007, 2010, 2011.
 - International Conference on Intelligent Robots and Systems (IROS), 2005, 2006, 2007, 2008, 2010.
 - International Conference on Advanced Video and Signal-Based Surveillance (AVSS), 2011.
- Memberships: IEEE Computer Society, ACM computing Membership

Invited Talks, Presentations

- Invited talk, “Big Data Computer Vision”, at SUNY Albany, Albany, NY, USA, 2013.
- Invited talk, “A Videography Analysis Framework for Video Retrieval and Summarization”, at IBM Research, Yorktown Heights, NY, USA, 2013.
- Invited talk, “Trends in Computer Vision”, at Pohang University of Science and Technology, South Korea, 2012.
- Invited talk, “Computer Vision for Internet Video Search”, at Bard College, NY, USA. 2012.
- Invited talk, “Endowing Switching Linear Dynamic Systems with Higher-order Temporal Structure”, jointly at Disney Research and Carnegie Mellon University, Pittsburgh, USA, 2009.
- Invited talk, “Switching Linear Dynamic Systems with Higher-order Temporal Structure”, at Computer Science and Engineering department at Seoul National University, South Korea, Jan 15th, 2009.
- Invited talk, “On-line Learning of the Traversability of Unstructured Terrain for Outdoor Robot Navigation”, Snowbird Learning Workshop, at Snowbird Learning Workshop, Snowbird, Utah, U.S.A. 2006.
- Invited talk, “Automatic Acquisition of 4D Urban Models and Proactive Auditory Service for Enhanced User Experience”, at KAIST, Korea, 2006.

Patents

- **Systems and Methods for Retrieving Causal Sets of Events from Unstructured Signals**, Jim Rehg, K. Prabhakar, Sangmin Oh, Ping Wang, Gregory Abowd, US Patent 20120301105
- **Fast Landmark Detection using Regression methods**, Cha Zhang, Paul Viola, and Sang Min Oh, US Patent 20080187213.

Software Skills

- Programming languages: Python, C++/C/C#, Matlab, Java, OCaml
- Version control tools: SVN, Git, CVS
- Platforms: Mac, Windows, Linux

Status

South Korea Citizenship. U.S. Permanent Resident.