

Jason Joseph Corso

Prepared on April 26, 2025

CONTACT	PHONE: 734-647-8833 EMAIL: jjcorso@umich.edu WEB: http://web.eecs.umich.edu/~jjcorso	PHONE: 734-531-9349 EMAIL: jason@voxel51.com WEB: https://voxel51.com
BACKGROUND	DATE OF BIRTH: 6 August 1978 CITIZENSHIP: USA	PLACE OF BIRTH: New York, USA SECURITY CLEARANCE: On Request
APPOINTMENTS	Professor (with tenure) at the University of Michigan Robotics (founded 7/2022) Electrical Engineering and Computer Science <i>Adjunct from 1/2021 to 7/2022</i> <i>Associate Professor (with tenure) from 8/2014 to 8/2019</i>	ANN ARBOR, MI 7/2022 - PRESENT 8/2019 - PRESENT
	Co-Founder at Voxel51, Inc. Chief Science Officer Chief Executive Officer Voxel51 is a Bessemer backed Series B startup company with 45+ employees building the best in class Open Source AI developer infrastructure, FiftyOne, with more than \$45M in venture funding.	ANN ARBOR, MI 7/2023 – PRESENT 12/2016 – 7/2023
	Advisory Board Member to the Department of Computer Science at Johns Hopkins University <i>Full list of appointments begins on page 33</i>	BALTIMORE, MD 7/2024 - PRESENT
RESEARCH FOCUS	Cognitive systems and their entangled challenges in vision, language, robotics, learning, AI, and the physical world, with applications in the health sciences, autonomy, and everyday life.	
RESEARCH AREAS	Computer Vision AI/Machine Learning	Robotics Human Interaction Natural Language Software Systems
EDUCATION	University of California, Los Angeles Post-Doc in Neuroscience, Radiology, and Statistics ADVISORS: Dr. Alan Yuille and Dr. Arthur Toga	LOS ANGELES, CA 2006-2007
	The Johns Hopkins University Ph.D. in Computer Science ADVISOR: Dr. Gregory D. Hager DISSERTATION TITLE: <i>“Techniques for Vision-Based Human-Computer Interaction”</i> <i>Full list of education begins on page 32</i>	BALTIMORE, MD 6/2006
DISTINCTIONS	EECS Outstanding Achievement Award <i>University of Michigan, Department of Electrical Engineering and Computer Science</i>	2018
	Best Associate Editor Award for ICRA	2016
	Google Faculty Research Award <i>Computational PB&J: Learning from Instructional Video</i>	2015
	Best Paper Award at ECDM 2012 for <i>Efficient Max-Margin Metric Learning</i>	2012
	SUNY at Buffalo Young Investigator Award	2011
	Army Research Office Young Investigator Award <i>Guidance By Semantics–High-Level Visual Inference to Improve Vision-based Mobile Robot Localization</i>	2010
	National Science Foundation CAREER Award <i>Generalized Image Understanding with Probabilistic Ontologies and Dynamic Adaptive Graph Hierarchies</i>	2009
	DARPA Computer Science Study Group <i>A distinction awarded to junior faculty for revolutionary activities in defense-relevant research</i>	2009

SUNY at Buffalo STOR Visionary Innovator <i>Awarded for licensing technology to industry.</i>	2009
Best New Development <i>UCLA Laboratory of Neuroimaging, CCB AHM Segmentation Contest</i>	2006
Link Foundation Fellowship in Advanced Simulation and Training	2003
James D Rozics Computer Science Medal - Loyola College in Maryland <i>Awarded to the computer science student ranked first upon graduation</i>	2000
Upsilon Pi Epsilon Scholarship, Computer Science Honors Society	1998
Hauber Summer Science Research Fellowship - Loyola College in Maryland	1998

PUBLICATIONS

Google Scholar (link): h-index 50 with 19K+ citations.

Publications are grouped by type into Journal, Peer-reviewed Conference, Other, Patents, and Reports; in-review publications are included with the Reports on page 11.

JOURNAL ARTICLES

Qualifiers from date of publication added where known: Impact Factor (IF), h5-Index (h5) provided by Google Scholar.

- J33. N. Louis, L. Zhou, S. J. Yule, R. D. Dias, M. Manojlovich, F. D. Pagani, D. S. Likosky, and **J. J. Corso**. Temporally guided articulated hand pose tracking in surgical videos. *International Journal of Computer Aided Radiology and Surgery*, 18(1):117–125, 2023.
IF: 2.3; h5: 51 (**IJCARS**)
- J32. M. R. Ganesh, Blanchard, **J. J. Corso**, and S. Y. Sekeh. Slimming neural networks using adaptive connectivity scores. *IEEE Transactions on Neural Networks and Learning Systems*, 35(3):3794–3808, 2022. (arXiv: 2006.12463)
IF: 11.8; h5: 145 (**TNNLS**)
- J31. M. R. Mathis, S. Yule, X. Wu, R. D. Dias, A. M. Janda, S. L. Krein, M. Manojlovich, M. D. Caldwell, K. Stakich-Alpirez, M. Zhang, **J. J. Corso**, N. Louis, T. Xu, J. Wolverson, F. D. Pagani, and D. Likosky. The impact of team familiarity on intra and postoperative cardiac surgical outcomes. *Surgery*, 170(4):1031–1038, 2021.
IF: 3.982; h5: 56 (**Surgery**)
- J30. D. Likosky, S. J. Yule, M. R. Mathis, R. D. Dias, **J. J. Corso**, M. Zhang, S. L. Krein, M. D. Caldwell, N. Louis, A. M. Janda, N. J. Shah, F. D. Pagani, K. Stakich-Alpirez, and M. M. Manojlovich. Novel assessments of technical and nontechnical cardiac surgery quality: Protocol for a mixed methods study. *JMIR Research Protocols*, 10(1), 2021.
IF: 1.25; h5: 32 (**JMIR**)
- J29. M. Mathis, S. Yule, X. Wu, R. D. Dias, A. M. Janda, S. L. Krein, M. Manojlovich, M. D. Caldwell, K. Stakich-Alpirez, M. Zhang, **J. J. Corso**, N. Louis, T. Xu, J. Wolverson, F. D. Pagani, and D. S. Likosky. Video assessment cardiac surgery quality. *Surgery*, 2021.
IF: 3.982; h5: 56 (**Surgery**)
- J28. R. Szeto, X. Sun, K. Lu, and **J. J. Corso**. A Temporally-Aware Interpolation Network for Video Frame Inpainting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(5), 2020.
IF: 17.730; h5: 127 (**IEEE TPAMI**)
- J27. Y. Yan, C. Xu, D. Cai, and **J. J. Corso**. A weakly supervised multi-task ranking framework for actor-action semantic segmentation. *International Journal of Computer Vision*, 128:1414–1432, 2020.
IF: 11.541; h5: 66 (**IJCV**)
- J26. S. Kumar, V. Dhiman, P. Koch, and **J. J. Corso**. Learning compositional sparse bimodal models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(5):1032–1044, 2018.
IF: 8.329; h5: 114 (**IEEE TPAMI**)
- J25. D. Sarikaya, **J. J. Corso**, and K. A. Guru. Detection and localization of robotic tools in robot-assisted surgery videos using deep neural networks for region proposal and detection. *IEEE Transactions on Medical Imaging*, 36(7):1542–1549, 2017.
IF: 3.76; h5: 59 (**IEEE TMI**)
- J24. T. Han, H. Yao, C. Xu, X. Sun, Y. Zhang, and **J. J. Corso**. Dancelets mining for video recommendation based on dance styles. *IEEE Transactions on Multimedia*, 19(4), 2017.
IF: 2.536; h5: 51 (**IEEE TMM**)
- J23. C. Chen and **J. J. Corso**. Joint occlusion boundary detection and figure/ground assignment by extracting common-fate fragments in a back-projection scheme. *Pattern Recognition*, 64:15–28, 2017.
IF: 3.399; h5: 67 (**PR**)

- J22. C. Xiong, D. M. Johnson, and **J. J. Corso**. Active clustering with model-based uncertainty reduction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(1):5–17, 2017. Original Version: ArXiv 1402.1783. *IF: 5.694; h5: 108 (IEEE TPAMI)*
- J21. C. Xu and **J. J. Corso**. LIBSVX: A supervoxel library and benchmark for early video processing. *International Journal of Computer Vision*, 119:272–290, 2016. *IF: 4.270; h5: 65 (IJCV)*
- J20. D. M. Johnson, C. Xiong, and **J. J. Corso**. Semi-supervised nonlinear distance metric learning via forests of max-margin cluster hierarchies. *IEEE Transactions on Knowledge and Data Engineering*, 28(4):1035–1046, 2016. *IF: 2.067; h5: 65 (IEEE TKDE)*
- J19. B. H. Menze, A. Jakab, S. Bauer, J. Kalpathy-Cramer, K. Farahani, J. Kirby, Y. Burren, N. Porz, J. Slotboom, R. Wiest, L. Lanczi, R. Gerstner, M.-A. Weber, T. Arbel, B. B. Avants, N. Ayache, P. Buendia, D. L. Collins, N. Cordier, **J. J. Corso**, A. Criminisi, T. Das, H. Delingette, C. Demiralp, C. R. Durst, M. Dojat, S. Doyle, J. Festa, F. Forbes, E. Geremia, B. Glocker, P. Golland, X. Guo, A. Hamamci, K. M. Iftekharuddin, R. Jena, N. M. John, E. Konukoglu, D. Lashkari, J. A. Mariz, R. Meier, S. Pereira, D. Precup, S. J. Price, T. Riklin Raviv, S. Reza, M. Ryan, D. Sarikaya, L. Schwartz, H.-C. Shin, J. Shotton, C. A. Silva, N. Sousa, N. K. Subbanna, G. Szekely, T. J. Taylor, O. M. Thomas, N. J. Tustison, G. Unal, F. Vasseur, M. Wintermark, D. H. Ye, L. Zhao, B. Zhao, D. Zikic, M. Prastawa, M. Reyes, and K. Van Leemput. The multimodal brain tumor image segmentation benchmark (brats). *IEEE Transactions on Medical Imaging*, 34(10):1993–2024, 2015. *IF: 3.39; h5: 63 (IEEE TMI)*
- J18. S. Kumar, N. Ahmidi, G. Hager, P. Singhal, **J. J. Corso**, and V. Krovi. Surgical performance assessment. *ASME Dynamics Systems and Control Magazine*, 3(3):7–10, 2015. **(ASME)**
- J17. C. Xu, R. F. Doell, S. J. Hanson, C. Hanson, and **J. J. Corso**. A study of actor and action semantic retention in video supervoxel segmentation. *International Journal of Semantic Computing*, 2014. Selected as a Best Paper from ICSC; an earlier version appeared as arXiv:1311.3318. *h5: 13 (IJSC)*
- J16. P. Agarwal, S. Kumar, J. Ryde, **J. J. Corso**, and V. N. Krovi. Estimating dynamics on-the-fly using monocular video for vision-based robotics. *IEEE/ASME Transactions on Mechatronics*, 19(4):1412–1423, 2014. *IF: 3.135; h5: 37 (IEEE/ASME TMECH)*
- J15. W. Wu, A. Y. C. Chen, L. Zhao, and **J. J. Corso**. Brain tumor detection and segmentation in a CRF (conditional random fields) framework with pixel-pairwise affinity and superpixel-level features. *International Journal of Computer Aided Radiology and Surgery*, 9(2):241–253, 2014. *IF: 1.364; h5: 17 (IJCARs)*
- J14. S. Oh, S. McCloskey, I. Kim, A. Vahdat, K. Cannons, H. Hajimirsadeghi, G. Mori, A. G. A. Perera, M. Pandey, and **J. J. Corso**. Multimedia event detection with multimodal feature fusion and temporal concept localization. *Machine Vision and Applications*, 25:49–69, 2014. *IF: 1.009; h5: 21 (MVA)*
- J13. J. A. Delmerico, P. David, and **J. J. Corso**. Building facade detection, segmentation, and parameter estimation for mobile robot stereo vision. *Image and Vision Computing*, 31(11):841–852, 2013. *IF: 1.952; h5: 42 (IVC)*
- J12. Y. Miao and **J. J. Corso**. Hamiltonian streamline guided feature extraction with application to face detection. *Journal of Neurocomputing*, 120:226–234, 2013. Early version appears as arXiv.org tech report 1108.3525v1. *IF: 1.58; h5: 44 (NEUCOM)*
- J11. **J. J. Corso**. Toward parts-based scene understanding with pixel-support parts-sparse pictorial structures. *Pattern Recognition Letters: Special Issue on Scene Understanding and Behavior Analysis*, 34(7):762–769, 2013. Early version appears as arXiv.org tech report 1108.4079v1. *IF: 1.034; h5: 43 (PRL)*
- J10. R. S. Alomari, **J. J. Corso**, V. Chaudhary, and G. Dhillon. Toward a clinical lumbar CAD: Herniation diagnosis. *International Journal of Computer Aided Radiology and Surgery*, 6(1):119–126, 2011. *IF: 1.364; h5: 17 (IJCARs)*
- J9. R. S. Alomari, **J. J. Corso**, and V. Chaudhary. Labeling of lumbar discs using both pixel- and object-level features with a two-level probabilistic model. *IEEE Transactions on Medical Imaging*, 30(1):1–10, 2011. *IF: 4.027; h5: 48 (IEEE TMI)*
- J8. P. B. Noël, A. Walczak, J. Xu, **J. J. Corso**, K. R. Hoffmann, and S. Schafer. GPU-based cone beam computed tomography. *Computer Methods and Programs in Biomedicine*, 98(3):271–277, 2010. *IF: 1.589; h5: 29 (CMPB)*

- J7. R. S. Alomari, **J. J. Corso**, V. Chaudhary, and G. Dhillon. Computer-aided diagnosis of lumbar disc pathology from clinical lower spine MRI. *International Journal of Computer Aided Radiology and Surgery*, 5(3):287–293, 2010. IF: 1.364; h5: 17 (**IJCARS**)
- J6. **J. J. Corso** and G. D. Hager. Image Description with Features that Summarize. *Computer Vision and Image Understanding*, 113:446–458, 2009. IF: 2.202; h5: 39 (**CVIU**)
- J5. **J. J. Corso**, G. Ye, D. Burschka, and G. D. Hager. A Practical Paradigm and Platform for Video-Based Human-Computer Interaction. *IEEE Computer*, 42(5):48–55, 2008. IF: 2.111; h5: 49 (**IEEE Computer**)
- J4. **J. J. Corso**, E. Sharon, S. Dube, S. El-Saden, U. Sinha, and A. Yuille. Efficient Multilevel Brain Tumor Segmentation with Integrated Bayesian Model Classification. *IEEE Transactions on Medical Imaging*, 27(5):629–640, 2008. IF: 4.027; h5: 48 (**IEEE TMI**)
- J3. **J. J. Corso**, G. Ye, and G. D. Hager. Analysis of Composite Gestures with a Coherent Probabilistic Graphical Model. *Virtual Reality*, 8(4):242–252, 2005. IF: 0.341; h5: 13 (**VR**)
- J2. D. Burschka, **J. J. Corso**, M. Dewan, W. Lau, M. Li, H. Lin, P. Marayong, N. Ramey, G. D. Hager, B. Hoffman, D. Larkin, and C. Hasser. Navigating Inner Space: 3-D Assistance for Minimally Invasive Surgery. *Robotics and Autonomous System*, 2005. IF: 1.615; h5: 37 (**RAS**)
- J1. G. Ye, **J. J. Corso**, D. Burschka, and G. D. Hager. VICs: A Modular HCI Framework Using Spatio-Temporal Dynamics. *Machine Vision and Applications*, 16(1):13–20, 2004. IF: 1.009; h5: 21 (**MVA**)

CONFERENCE PROCEEDINGS (PEER-REVIEWED)

Qualifiers added where known: Acceptance Rate* (AR) with * indicating historical, not year-specific rate, h5-Index (h5) provided by Google Scholar (at time of publication).

- 2024 C129. N. Louis, M. Khoshlessan, and **J. J. Corso**. Measuring physical plausibility of 3D human poses using physical simulation. In *Proceedings of the British Machine Vision Conference, 2024*. AR: 25.8%; h5: 65 (**BMVC**)
- 2023 C128. J. Krantz, S. Banerjee, W. Zhu, **J. J. Corso**, P. Anderson, S. Lee, and J. Thomason. Iterative vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023*. (arXiv: 2210.03087) AR: 25.8%; h5: 440 (**CVPR**)
- C127. S. J. Lemmer, A. Guo, and **J. J. Corso**. Human-centered deferred inference: Measuring user interactions and setting deferral criteria for human-ai teams. In *Proceedings of the International Conference on Intelligent User Interfaces, 2023*. AR: 27%; h5: 52 (**IUI**)
- C126. S. J. Lemmer and **J. J. Corso**. Evaluating and improving interactions with hazy oracles. In *Proceedings of the AAAI Conference on Artificial Intelligence, 2023*. (arXiv: 2110.10206) AR: 19.6%; h5: 220 (**AAAI**)
- C125. M. R. Ganesh, S. Y. Sekeh, and **J. J. Corso**. Can deep networks be highly performance, efficient and robust simultaneously. In *Proceedings of the British Machine Vision Conference, 2023*. AR: 25.3%; h5: 65 (**BMVC**)
- 2022 C124. R. Szeto and **J. J. Corso**. The DEVIL is in the details: A diagnostic evaluation benchmark for video inpainting. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022*. (arXiv: 2105.05332) AR: 24.9%; h5: 440 (**CVPR**)
- C123. N. Louis, **J. J. Corso**, T. N. Templin, T. D. Eliason, and D. P. Nicolella. Learning to estimate external forces of human motion in video. In *Proceedings of ACM International Conference on Multimedia, 2022*. (arXiv: 2207.05845) AR: 27.9%; h5: 101 (**ACMMM**)
- 2021 C122. S. J. Lemmer and **J. J. Corso**. Ground-truth or DAER: Selective re-query of secondary information. In *Proceedings of IEEE International Conference on Computer Vision, 2021*. AR: 25.9%; h5: 184 (**ICCV**)
- C121. S. Banerjee, J. Thomason, and **J. J. Corso**. The RobotSLANG benchmark: Dialog-guided robot localization and navigation. In *Proceedings of Conference on Robot Learning, 2021*. AR: 34.7%; h5: 49 (**CoRL**)
- C120. B. A. Griffin and **J. J. Corso**. Depth from camera motion and object detection. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021*. AR: 22%; h5: 356 (**CVPR**)

- C119. S. J. Lemmer, J. Y. Song, and **J. J. Corso**. Crowdsourcing more effective initializations for single-target trackers through automatic re-querying. In *Proceedings of ACM Conference on Computer Human Interaction*, 2021. AR: 26.3%; h5: 101 (**CHI**)
- C118. G. Liu, H. Tang, H. Latapie, **J. J. Corso**, and Yan Yan. Cross-view exocentric to egocentric video synthesis. In *Proceedings of ACM International Conference on Multimedia*, 2021. AR: 27.9%; h5: 65 (**ACMMM**)
- C117. R. Szeto, M. El-Khamy, J. Lee, and **J. J. Corso**. HyperCon: Image-to-video model transfer for video-to-video translation tasks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021. AR*: 35%; h5: 62 (**WACV**)
- C116. K. Min and **J. J. Corso**. Integrating human gaze into attention for edocentric activity recognition. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021. AR*: 35%; h5: 62 (**WACV**)
- 2020 C115. L. Zhou, H. Palangi, L. Zhang, H. Hu, **J. J. Corso**, and J. Gao. Unified vision-language pre-training for image captioning and vqa. In *Proceedings of AAAI Conference on Artificial Intelligence*, 2020. AR: 20.6%; h5: 95 (**AAAI**)
- C114. B. Griffin and **J. J. Corso**. Learning object depth from camera motion and video object segmentation. In *Proceedings of European Conference on Computer Vision*, 2020. AR: 26%; h5: 144 (**ECCV**)
- C113. K. Min and **J. J. Corso**. Adversarial background-aware loss for weakly-supervised temporal activity localization. In *Proceedings of European Conference on Computer Vision*, 2020. AR: 26%; h5: 144 (**ECCV**)
- C112. M. El Banani, **J. J. Corso**, and D. Fouhey. Novel object viewpoint estimation through reconstruction alignment. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2020. AR: 22%; h5: 299 (**CVPR**)
- C111. M. Ravi Ganesh and **J. J. Corso**. Rethinking curriculum learning with incremental labels and adaptive compensation. In *Proceedings of the British Machine Vision Conference*, 2020. AR: 29%; h5: 57 (**BMVC**)
- C110. V. Florence, **J. J. Corso**, and B. Griffin. Robot-supervised learning for object segmentation. In *Proceedings of IEEE International Conference on Robotics and Automation*, 2020. AR: 42%; h5: 82 (**ICRA**)
- C109. B. Griffin, V. Florence, and **J. J. Corso**. Video object segmentation-based visual servo control and object depth estimation on a mobile robot. In *Proceedings of IEEE Winter Conference on Applications of Computer Vision*, 2020. AR*: 35%; h5: 46 (**WACV**)
- C108. M. R. Ganesh, **J. J. Corso**, and S. Y. Sekeh. MINT: Deep network compression via mutual information-based neuron trimming. In *Proceedings of IEEE Conference on Pattern Recognition*, 2020. AR*: 50%; h5: 43 (**ICPR**)
- 2019 C107. K. Min and **J. J. Corso**. TASED-net: Temporally-aggregating spatial encoder-decoder network for video saliency detection. In *Proceedings of IEEE International Conference on Computer Vision*, 2019. AR: 25%; h5: 129 (**ICCV**)
- C106. B. Griffin and **J. J. Corso**. BubbleNets: Learning to select the guidance frame in video object segmentation by deep sorting frames. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2019. AR: 5.6% (Oral; Best Paper Finalist); h5: 188 (**CVPR**)
- C105. L. Zhou, Y. Kalantidis, X. Chen, **J. J. Corso**, and M. Rohrbach. Grounded video description. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2019. AR: 5.6% (Oral); h5: 188 (**CVPR**)
- C104. H. Tang, D. Xu, Y. Yan, Y. Wang, **J. J. Corso**, and N. Sebe. Multi-channel attention selection GAN with cascaded semantic guidance for cross-view image translation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2019. AR: 5.6% (Oral); h5: 188 (**CVPR**)
- C103. H. Huang, L. Zhou, W. Zhang, **J. J. Corso**, and C. Xu. Dynamic graph modules for modeling object-object interactions in activity recognition. In *Proceedings of the British Machine Vision Conference*, 2019. AR: 28%; h5: 46 (**BMVC**)
- C102. H. Tang, X. Chen, W. Wang, D. Xu, **J. J. Corso**, N. Sebe, and Y. Yan. Attribute-guided sketch generation. In *Proceedings of IEEE Conference on Automatic Face and Gesture Recognition*, 2019. AR*: 30%; h5: 36 (**FG**)

- C101. J. Y. Song, S. J. Lemmer, M. X. Liu, S. Yan, J. Kim, **J. J. Corso**, and W. S. Lasecki. Popup: Reconstructing 3d video using particle filtering to aggregate crowd responses. In *Proceedings of ACM International Conference on Intelligent User Interfaces*, 2019. AR: 25%; h5: 27 (**IUI**)
- C100. B. Griffin and **J. J. Corso**. Tukey-inspired video object segmentation. In *Proceedings of IEEE Winter Conference on Applications of Computer Vision*, 2019. AR: 37%; h5: 38 (**WACV**)
- 2018 C99. L. Zhou, C. Xu, and **J. J. Corso**. Towards automatic learning of procedures from web instructional videos. In *Proceedings of AAAI Conference on Artificial Intelligence*, 2018. AR: 24.6%; h5: 56 (**AAAI**)
- C98. L. Zhou, Y. Zhou, **J. J. Corso**, R. Socher, and C. Xiong. End-to-end dense video captioning with masked transformer. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2018. AR: 29.5%; h5: 158 (**CVPR**)
- C97. X. Sun, R. Szeto, and **J. J. Corso**. A Temporally-Aware Interpolation Network for Video Frame Inpainting. In *Proceedings of Asian Conference on Computer Vision (ACCV)*, 2018. AR: 28%; h5: 30 (**ACCV**)
- C96. L. Zhou, N. Louis, and **J. J. Corso**. Weakly-supervised video object grounding from text by loss weighting and object interaction. In *Proceedings of British Machine Vision Conference*, 2018. AR: 30%; h5: 42 (**BMCV**)
- C95. K. R. Keane and **J. J. Corso**. The wrong tool for inference — a critical view of gaussian graphical models. In *Proceedings of the 7th International Conference on Pattern Recognition Applications and Methods*, 2018. AR*: 50%; h5: 12 (**ICPRAM**)
- 2017 C94. R. Szeto and **J. J. Corso**. Click-here: Human-localized keypoints as guidance for viewpoint estimation. In *Proceedings of IEEE International Conference on Computer Vision*, 2017. AR: 28.9%; h5: 89 (**ICCV**)
- C93. Y. Yan, C. Xu, D. Cai, and **J. J. Corso**. Weakly supervised actor-action segmentation via robust multi-task ranking. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2017. AR: 28%; h5: 140 (**CVPR**)
- C92. L. Zhou, C. Xu, P. Koch, and **J. J. Corso**. Watch what you just said: Image captioning with text-conditional attention. In *Proceedings of the Thematic Workshops of ACM Multimedia*, 2017. (**ACMMMWW**)
- 2016 C91. V. Dhiman, Q.-H. Tran, **J. J. Corso**, and M. Chandraker. A continuous occlusion model for road scene understanding. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2016. AR*: 25%; h5: 128 (**CVPR**)
- C90. C. Xu and **J. J. Corso**. Actor-action semantic segmentation with grouping-process models. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2016. AR*: 25%; h5: 128 (**CVPR**)
- C89. W. Ding, M. R. Ganesh, R. N. Severinghaus, **J. J. Corso**, and D. Panagou. Real-time model predictive control for keeping a quadrotor vehicle on the camera field-of-view of a ground robot. In *Proceedings of the American Control Conference*, 2016. AR*: 50%; h5: 43 (**ACC**)
- C88. L. Zhao, D. Chan, Z. Chen, X. Wang, N. Paliwal, J. Xiang, H. Meng, **J. J. Corso**, and J. Xu. Rapid virtual stenting for intracranial aneurysms. In *Proceedings of SPIE Conference on Medical Imaging*, 2016. h5: 56 (**SPIE**)
- 2015 C87. W. Chen and **J. J. Corso**. Action detection by implicit intentional motion clustering. In *Proceedings of IEEE International Conference on Computer Vision*, 2015. AR: 30.3%; h5: 68 (**ICCV**)
- C86. C. Xu, S.-H. Hsieh, C. Xiong, and **J. J. Corso**. Can humans fly? Action understanding with multiple classes of actors. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2015. AR*: 25%; h5: 118 (**CVPR**)
- C85. J. Lu, R. Xu, and **J. J. Corso**. Human action segmentation with hierarchical supervoxel consistency. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2015. AR*: 25%; h5: 118 (**CVPR**)
- C84. R. Xu, C. Xiong, W. Chen, and **J. J. Corso**. Jointly modeling deep video and compositional text to bridge vision and language in a unified framework. In *Proceedings of AAAI Conference on Artificial Intelligence*, 2015. AR: 27%; h5: 37 (**AAAI**)

2014

- C83. S. Kumar, V. Dhiman, and **J. J. Corso**. Learning compositional sparse models of bimodal percepts. In *Proceedings of AAAI Conference on Artificial Intelligence*, 2014. AR: 28%; h5: 41 (AAAI)
- C82. C. Xiong, S. McCloskey, and **J. J. Corso**. Latent domains for visual domain adaptation. In *Proceedings of AAAI Conference on Artificial Intelligence*, 2014. AR: 28%; h5: 41 (AAAI)
- C81. W. Chen, C. Xiong, R. Xu, and **J. J. Corso**. Actionness ranking with lattice conditional ordinal random fields. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2014. AR: 29%; h5: 106 (CVPR)
- C80. A. Barbu, D. Barrett, W. Chen, N. Siddharth, C. Xiong, **J. J. Corso**, C. D. Fellbaum, C. Hanson, S. J. Hanson, S. Hélie, E. Malaia, B. A. Pearlmutter, J. M. Siskind, T. M. Talavage, and R. B. Wilbur. Seeing is worse than believing: Reading people’s minds better than computer-vision methods recognize actions. In *Proceedings of European Conference on Computer Vision*, 2014. AR: 24%; h5: 59 (ECCV)
- C79. V. Dhiman, A. Kundu, F. Dellaert, and **J. J. Corso**. Modern MAP inference methods for accurate and faster occupancy grid mapping on higher order factor graphs. In *Proceedings of International Conference on Robotics and Automation*, 2014. AR: 48%; h5: 55 (ICRA)
- C78. S. Kumar, M. S. Narayanan, P. Singhal, **J. J. Corso**, and V. Krovi. Surgical tool attributes from monocular video. In *Proceedings of IEEE International Conference on Robotics and Automation*, 2014. AR: 48%; h5: 55 (ICRA)
- C77. C. Xiong, W. Chen, G. Chen, D. Johnson, and **J. J. Corso**. Adaptive quantization: An information-based approach to learning binary codes. In *Proceedings of SIAM International Conference on Data Mining*, 2014. AR*: 15%; h5: 33 (SDM)
- C76. G. A. Gross, D. R. Schlegel, **J. J. Corso**, J. Llinas, R. Nagi, and S. C. Shapiro. Systemic test and evaluation of a hard+soft information fusion framework: Challenges and current approaches. In *Proceedings of International Conference on Information Fusion*, 2014. (FUSION)

2013

- C75. C. Xu, S. Whitt, and **J. J. Corso**. Flattening supervoxel hierarchies by the uniform entropy slice. In *Proceedings of the IEEE International Conference on Computer Vision*, 2013. AR*: 20%; h5: 60 (ICCV)
- C74. C. Xu, R. F. Doell, S. J. Hanson, C. Hanson, and **J. J. Corso**. Are actor and action semantics retained in video supervoxel segmentation? In *Proceedings of IEEE International Conference on Semantic Computing*, 2013. AR*: 30% (ICSC)
- C73. V. Dhiman, J. Ryde, and **J. J. Corso**. Mutual localization: Two camera relative 6-dof pose estimation from reciprocal fiducial observation. In *Proceedings of International Conference on Intelligent Robots and Systems*, 2013. AR*: 43%; h5: 40 (IROS)
- C72. S. Kumar, M. Narayanan, P. Singhal, **J. J. Corso**, and V. Krovi. Product of tracking experts for surgical tool visual tracking. In *IEEE Conference on Automation Science and Engineering*, 2013. h5: 13 (CASE)
- C71. L. Zhao, W. Wu, and **J. J. Corso**. Semi-automatic brain tumor segmentation by constrained MRFs using structural trajectories. In *Proceedings of Medical Image Computing and Computer Aided Intervention*, 2013. AR*: 30%; h5: 30 (MICCAI)
- C70. P. Das, C. Xu, R. F. Doell, and **J. J. Corso**. A thousand frames in just a few words: Lingual description of videos through latent topics and sparse object stitching. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2013. AR*: 25%; h5: 106 (CVPR)
- C69. D. M. Johnson, C. Xiong, J. Gao, and **J. J. Corso**. Comprehensive cross-hierarchy cluster agreement evaluation. In *Proceedings of AAAI Conference on Artificial Intelligence (Late-Breaking Papers Track)*, 2013. AR: 43%; h5: 41 (AAAI)
- C68. C. Xiong, D. M. Johnson, and **J. J. Corso**. Uncertainty reduction for active image clustering via a hybrid global-local uncertainty model. In *Proceedings of AAAI Conference on Artificial Intelligence (Late-Breaking Papers Track)*, 2013. AR: 43%; h5: 41 (AAAI)
- C67. N. Coffee, J. Gawley, C. W. Forstall, W. J. Scheirer, D. Johnson, **J. J. Corso**, and B. Parks. Modelling the interpretation of literary allusion with machine learning techniques. In *Proceedings of Digital Humanities*, 2013. (DH)
- C66. P. Das, R. K. Srihari, and **J. J. Corso**. Translating related words to videos and back through latent topics. In *Proceedings of Sixth ACM International Conference on Web Search and Data Mining*, 2013. AR: 19%; h5: 48 (WSDM)

- C65. J. A. Delmerico, D. Baran, P. David, J. Ryde, and **J. J. Corso**. Ascending stairway modeling from dense depth imagery for traversability analysis. In *Proceedings of IEEE International Conference on Robotics and Automation*, 2013. AR*: 45%; h5: 55 (**ICRA**)
- C64. R. S. Alomari, **J. J. Corso**, V. Chaudhary, and G. Dhillon. Lumbar spine disc herniation diagnosis with a joint shape model. In *Proceedings of Medical Image Computing and Computer Aided Intervention Workshop on Computational Spine Imaging*, 2013. (**MICCAIW**)
- 2012 C63. C. Xu, C. Xiong, and **J. J. Corso**. Streaming hierarchical video segmentation. In *Proceedings of European Conference on Computer Vision*, 2012. AR*: 26%; h5: 66 (**ECCV**)
- C62. J. Ryde and **J. J. Corso**. Fast voxel maps with counting bloom filters. In *Proceedings of International Conference on Intelligent Robots and Systems*, 2012. AR*: 43%; h5: 40 (**IROS**)
- C61. J. A. Delmerico, **J. J. Corso**, D. Baran, P. David, and J. Ryde. Ascending stairway modeling: A first step toward autonomous multi-floor exploration. In *Proceedings of IEEE/RSJ Intelligent Robots and Systems (Video Proceedings)*, 2012. AR*: 43%; h5: 40 (**IROS**)
- C60. C. Xiong, D. Johnson, R. Xu, and **J. J. Corso**. Random forests for metric learning with implicit pairwise position dependence. In *Proceedings of ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2012. AR*: 18%; h5: 67 (**KDD**)
- C59. S. Sadanand and **J. J. Corso**. Action bank: A high-level representation of activity in video. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2012. AR*: 25%; h5: 106 (**CVPR**)
- C58. C. Xu and **J. J. Corso**. Evaluation of super-voxel methods for early video processing. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2012. AR*: 25%; h5: 106 (**CVPR**)
- C57. P. Agarwal, S. Kumar, J. Ryde, **J. J. Corso**, and V. N. Krovi. Estimating human dynamics on-the-fly using monocular video for pose estimation. In *Proceedings of Robotics Science and Systems*, 2012. AR*: 20%; h5: 29 (**RSS**)
- C56. R. Xu, P. Agarwal, S. Kumar, V. N. Krovi, and **J. J. Corso**. Combining skeletal pose with local motion for human activity recognition. In *Proceedings of VII Conference on Articulated Motion and Deformable Objects*, 2012. h5: 9 (**AMDO**)
- C55. G. Chen, C. Xiong, and **J. J. Corso**. Dictionary transfer for image denoising via domain adaptation. In *Proceedings of IEEE International Conference on Image Processing*, 2012. AR*: 45%; h5: 34 (**ICIP**)
- C54. K. R. Keane and **J. J. Corso**. Maintaining prior distributions across evolving eigenspaces: An application to portfolio construction. In *Proceedings of 11th International Conference on Machine Learning and Applications*, 2012. AR*: 40%; h5: 12 (**ICMLA**)
- C53. C. Xiong and **J. J. Corso**. Coaction discovery: Segmentation of common actions across multiple videos. In *Proceedings of Multimedia Data Mining Workshop in Conjunction with the ACM SIGKDD Conference on Knowledge Discovery and Data Mining (MDMKDD)*, 2012. (**MDMKDD**)
- C52. C. Xiong, D. Johnson, and **J. J. Corso**. Spectral active clustering via purification of the k -nearest neighbor graph. In *Proceedings of European Conference on Data Mining*, 2012. AR: 14% (**ECDM**)
- C51. C. Xiong, D. Johnson, and **J. J. Corso**. Efficient max-margin metric learning. In *Proceedings of European Conference on Data Mining*, 2012. **Winner of Best Paper Award at ECDM 2012..** AR: 14% (**ECDM**)
- C50. K. R. Keane and **J. J. Corso**. Dynamically mixing dynamic linear models with applications in finance. In *Proceedings of International Conference on Pattern Recognition Applications and Methods*, 2012. (**ICPRAM**)
- C49. M. A. Bustamante and **J. J. Corso**. Using probabilistic ontologies for video exploration. In *Proceedings of the Eighteenth Americas Conference on Information Systems*, 2012. (**AMCIS**)
- C48. P. Agarwal, S. Kumar, J. Ryde, **J. J. Corso**, and V. N. Krovi. An optimization based framework for human pose estimation in monocular videos. In *Proceedings of International Symposium on Visual Computing*, 2012. (**ISVC**)
- 2011 C47. J. A. Delmerico, P. David, and **J. J. Corso**. Building facade detection, segmentation, and parameter estimation for mobile robot localization and guidance. In *Proceedings of International Conference on Intelligent Robots and Systems*, 2011. AR*: 43%; h5: 40 (**IROS**)

- C46. P. Agarwal, S. Kumar, **J. J. Corso**, and V. N. Krovi. Estimating dynamics on-the-fly using monocular video. In *Proceedings of 4th Annual Dynamic Systems and Control Conference*, 2011. (DSCC)
- C45. D. Gagneja, C. Xiong, and **J. J. Corso**. Towards a parts-based approach to sub-cortical brain structure parsing. In *Proceedings of SPIE Conference on Medical Imaging*, 2011. h5: 56 (SPIE)
- C44. A. Y. C. Chen and **J. J. Corso**. Temporally consistent multi-class video-object segmentation with the video graph-shifts algorithm. In *Proceedings of the 2011 IEEE Workshop on Motion and Video Computing*, 2011. (WMVC)
- C43. D. R. Schlegel, A. Y. C. Chen, C. Xiong, J. A. Delmerico, and **J. J. Corso**. AirTouch: Interacting with computer systems at a distance. In *Proceedings of IEEE Winter Vision Meetings: Workshop on Applications of Computer Vision (WACV)*, 2011. h5: 16 (WACV)
- 2010 C42. W. Ceusters, **J. J. Corso**, Y. Fu, M. Petropoulos, and V. Krovi. Introducing ontological realism for semi-supervised detection and annotation of operationally significant activity in surveillance videos. In *Proceedings of the 5th International Conference on Semantic Technologies for Intelligence, Defense and Security (STIDS)*, 2010. (STIDS)
- C41. R. S. Alomari, **J. J. Corso**, V. Chaudhary, and G. Dhillon. Lumbar disc herniation cad with a GVF-snake model. In *Proceedings of the 24th International Conference on Computer Aided Diagnosis and Surgery (CARS '10)*, 2010. (CARS)
- C40. M. R. Malgireddy, **J. J. Corso**, S. Setlur, V. Govindaraju, and D. Mandalapu. A framework for hand gesture recognition and spotting using sub-gesture modeling. In *Proceedings of the 20th International Conference on Pattern Recognition*, 2010. AR*: 55%; h5: 31 (ICPR)
- C39. Y. Tang, S. Srihari, H. Kasiviswanathan, and **J. J. Corso**. Footwear print retrieval system for real crime scene marks. In *Proceedings of International Workshop on Computational Forensics*, 2010. (IWCF)
- C38. A. Y. C. Chen and **J. J. Corso**. On the effects of normalization in adaptive MRF hierarchies. In *Proceedings of CompImage '10—Computational Modeling of Objects Presented in Images*, 2010.
- C37. R. S. Alomari, **J. J. Corso**, V. Chaudhary, and G. Dhillon. Automatic diagnosis of lumbar disc herniation using shape and appearance features from mri. In *Proceedings of SPIE Conference on Medical Imaging*, 2010. h5: 56 (SPIE)
- 2009 C36. R. Rodrigues, G. Schroeder, **J. J. Corso**, and V. Govindaraju. Unconstrained face recognition using MRF priors and manifold traversing. In *Proceedings of IEEE International Conference on Biometrics: Theory, Applications, Systems*, 2009. AR*: 43% h5: 25 (BTAS)
- C35. Y. Tao, L. Lu, M. Dewan, A. Y. C. Chen, **J. J. Corso**, J. Xuan, M. Salganicoff, and A. Krishnan. Multi-level ground glass nodule detection and segmentation in ct lung images. In *Proceedings of Medical Image Computing and Computer Aided Intervention (MICCAI)*, volume LNCS 5762, pages 715–723, 2009. AR*: 30%; h5: 30 (MICCAI)
- C34. T. J. Burns and **J. J. Corso**. Robust unsupervised segmentation of degraded document images with topic models. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2009. AR*: 25%; h5: 106 (CVPR)
- C33. R. S. Alomari, **J. J. Corso**, V. Chaudhary, and G. Dhillon. Desiccation diagnosis in lumbar discs from clinical mri with a probabilistic model. In *Proceedings of 2009 IEEE International Symposium on Biomedical Imaging*, 2009. AR*: 60%; h5: 27 (ISBI)
- C32. H. Z. Girgis, **J. J. Corso**, and D. Fischer. On-line hierarchy of general linear models for selecting and ranking the best predicted protein structures. In *Proceedings of IEEE Conference on Engineering in Medicine and Biology*, volume 1, pages 4949–4953, 2009. h5: 28 (EMBS)
- C31. R. S. Alomari, **J. J. Corso**, V. Chaudhary, and G. Dhillon. Abnormality detection in lumbar discs from clinical mr images with a probabilistic model. In *Proceedings of 23rd International Congress and Exhibition on Computer Assisted Radiology and Surgery (CARS 2009)*, 2009. (CARS)
- C30. P. Noël, J. Xu, K. R. Hoffmann, and **J. J. Corso**. Geometric tomography: a limited-view approach for computed tomography. In *Proceedings of the 25th Annual Symposium on Computational Geometry*, 2009. AR*: 28% (ASCG)
- C29. C. S. Hoeflich and **J. J. Corso**. Segmentation of 2D Gel Electrophoresis Spots using a Markov Random Field. In *Proceedings of SPIE Conference on Medical Imaging*, 2009. h5: 56 (SPIE)
- C28. S. Seshamani, M. D. Smith, **J. J. Corso**, M. O. Filipovich, A. Natarajan, and G. D. Hager. Direct Global Adjustment Methods for Endoscopic Mosaicking. In *Proceedings of SPIE Conference on Medical Imaging*, 2009. h5: 56 (SPIE)

- C27. P. B. Noël, **J. J. Corso**, J. Xu, K. R. Hoffmann, S. Schafer, and A. Walczak. Reconstruction from a Flexible Number of Projections in Cone-Beam Computed Tomography via Active Shape Models. In *Proceedings of SPIE Conference on Medical Imaging*, 2009. h5: 56 (SPIE)
- C26. P. B. Noël, J. Xu, K. R. Hoffmann, **J. J. Corso**, S. Schafer, and A. Walczak. High Contrast Artifact Reduction in Cone Beam Computed Tomography by Using Geometric Techniques. In *Proceedings of SPIE Conference on Medical Imaging*, 2009. h5: 56 (SPIE)
- 2008 C25. A. Y. C. Chen, **J. J. Corso**, and L. Wang. HOPS: Efficient region labeling using higher order proxy neighborhoods. In *Proceedings of International Conference on Pattern Recognition*, 2008. AR*: 55%; h5: 31 (ICPR)
- C24. J. Li, S. Tulyakov, F. Farooq, **J. J. Corso**, and V. Govindaraju. Integrating minutiae based fingerprint matching with local mutual information. In *Proceedings of International Conference on Pattern Recognition*, 2008. AR*: 55%; h5: 31 (ICPR)
- C23. **J. J. Corso**, R. S. Alomari, and V. Chaudhary. Lumbar disc localization and labeling with a probabilistic model on both pixel and object features. In *Proceedings of Medical Image Computing and Computer Aided Intervention (MICCAI)*, volume LNCS 5241 Part 1, pages 202–210, 2008. AR*: 30%; h5: 30 (MICCAI)
- C22. I. Nwogu and **J. J. Corso**. Exploratory identification of image-based bio-markers for solid mass pulmonary tumors. In *Proceedings of Medical Image Computing and Computer Aided Intervention (MICCAI)*, volume LNCS 5241, Part 1, pages 612–619, 2008. AR*: 30%; h5: 30 (MICCAI)
- C21. **J. J. Corso**. Discriminative Modeling by Boosting on Multilevel Aggregates. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2008. AR*: 25%; h5: 106 (CVPR)
- C20. **J. J. Corso**, A. Yuille, and Z. Tu. Graph-Shifts: Natural Image Labeling by Dynamic Hierarchical Computing. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2008. AR*: 25%; h5: 106 (CVPR)
- C19. I. Nwogu and **J. J. Corso**. (BP)²: Beyond Pairwise Belief Propagation, Labeling by Approximating Kikuchi Free Energies. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2008. AR*: 25%; h5: 106 (CVPR)
- C18. **J. J. Corso**, Z. Tu, and A. Yuille. MRF Labeling with a Graph-Shifts Algorithm. In *Proceedings of International Workshop on Combinatorial Image Analysis*, volume LNCS 4958, pages 172–184, 2008. (IWICIA)
- C17. I. Nwogu and **J. J. Corso**. Labeling irregular graphs with belief propagation. In *Proceedings of International Workshop on Combinatorial Image Analysis*, volume LNCS 4958, pages 295–305, 2008. (IWICIA)
- C16. P. B. Noël, A. Walczak, K. R. Hoffmann, J. Xu, **J. J. Corso**, and S. Schafer. Clinical Evaluation of GPU-Based Cone Beam Computed Tomography. In *Proceedings of High-Performance Medical Image Computing and Computer-Aided Intervention (HP-MICCAI)*, 2008.
- C15. S. Dube, **J. J. Corso**, A. Yuille, T. F. Cloughesy, S. El-Saden, and U. Sinha. Hierarchical Segmentation of Malignant Gliomas Via Integrated Contextual Filter Response. In *Proceedings of SPIE Conference on Medical Imaging*, 2008. h5: 56 (SPIE)
- 2007 C14. **J. J. Corso**, A. L. Yuille, N. L. Sicotte, and A. Toga. Detection and Segmentation of Pathological Structures by the Extended Graph-Shifts Algorithm. In *Proceedings of Medical Image Computing and Computer Aided Intervention (MICCAI)*, 2007. AR*: 30%; h5: 30 (MICCAI)
- C13. **J. J. Corso**, Z. Tu, A. Yuille, and A. W. Toga. Segmentation of Sub-Cortical Structures by the Graph-Shifts Algorithm. In *Proceedings of Information Processing in Medical Imaging (IPMI)*, volume LNCS 4584, pages 183–197, 2007. AR: 28%; h5: 16 (IPMI)
- 2006 C12. **J. J. Corso**, E. Sharon, and A. L. Yuille. Multilevel Segmentation and Integrated Bayesian Model Classification with an Application to Brain Tumor Segmentation. In *Proceedings of Medical Image Computing and Computer Aided Intervention (MICCAI)*, volume 2, pages 790–798, 2006. AR*: 30%; h5: 30 (MICCAI)
- 2005 C11. **J. J. Corso** and G. D. Hager. Coherent Regions for Concise and Stable Image Description. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 184–190, 2005. AR*: 25%; h5: 106 (CVPR)

- 2004
- C10. **J. J. Corso**, M. Dewan, and G. D. Hager. Image Segmentation Through Energy Minimization Based Subspace Fusion. In *Proceedings of 17th International Conference on Pattern Recognition (ICPR 2004)*, 2004. AR*: 55%; h5: 31 (ICPR)
- C9. W. W. Lau, N. A. Ramey, **J. J. Corso**, N. Thakor, and G. D. Hager. Stereo-Based Endoscopic Tracking of Cardiac Surface Deformation. In *Proceedings of Seventh International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2004. AR*: 30%; h5: 30 (MICCAI)
- C8. N. Ramey, **J. J. Corso**, W. W. Lau, D. Burschka, and G. D. Hager. Real Time 3D Surface Tracking and Its Applications. In *Proceedings of Workshop on Real-time 3D Sensors and Their Use (at CVPR 2004)*, 2004. (CVPRW)
- C7. G. Ye, **J. J. Corso**, and G. D. Hager. Gesture Recognition Using 3D Appearance and Motion Features. In *Proceedings of Workshop on Real-time Vision for Human-Computer Interaction (at CVPR 2004)*, 2004. (CVPRW)
- 2003
- C6. G. Ye, **J. J. Corso**, G. D. Hager, and A. M. Okamura. VisHap: Augmented Reality Combining Haptics and Vision. In *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, 2003. h5: 20 (SMC)
- C5. **J. J. Corso**, D. Burschka, and G. D. Hager. Direct Plane Tracking in Stereo Image for Mobile Navigation. In *Proceedings of International Conference on Robotics and Automation*, 2003. AR*: 45%; h5: 55 (ICRA)
- C4. G. Ye, **J. J. Corso**, D. Burschka, and G. D. Hager. VICs: A Modular Vision-Based HCI Framework. In *Proceedings of 3rd International Conference on Computer Vision Systems*, pages 257–267, 2003. AR: 40%; h5: 15 (ICVS)
- C3. **J. J. Corso**, D. Burschka, and G. D. Hager. The 4DT: Unencumbered HCI With VICs. In *Proceedings of CVPRHCI*, 2003. (CVPRW)
- 2002
- C2. J. Leven, **J. J. Corso**, J. D. Cohen, and S. Kumar. Interactive Visualization of Unstructured Grids Using Hierarchical 3D Textures. In *Proceedings of IEEE/SIGGRAPH Symposium on Volume Visualization and Graphics 2002*, pages 37–44, 2002. (VOLVIS)
- C1. **J. J. Corso**, J. Chhugani, and A. Okamura. Interactive Haptic Rendering of Deformable Surfaces Based on the Medial Axis Transform. In *Eurohaptics*, pages 92–98, 2002.

OTHER: BOOK CHAPTERS AND THESES

- O5. L. Zhao, Z. Peng, K. Finkler, A. Jerebko, **J. J. Corso**, and X. S. Zhou. Automatic collimation detection in digital radiographs with the directed hough transform and learning-based edge detection. In *Patch-Based Techniques in Medical Imaging (LNCS 9467)*, pages 71–78. Springer International Publishing, 2015.
- O4. S. Dube, **J. J. Corso**, T. F. Cloughesy, S. El-Saden, A. Yuille, and U. Sinha. *Data Mining Systems Analysis and Optimization in Biomedicine*, chapter Automated MR Image Processing and Analysis of Malignant Brain Tumors: Enabling Technology for Data Mining. American Institute of Physics, 2007.
- O3. **J. J. Corso**. *Techniques for Vision-Based Human-Computer Interaction*. PhD thesis, The Johns Hopkins University, 2005.
- O2. G. Ye, **J. J. Corso**, and G. D. Hager. *Real-Time Vision for Human-Computer Interaction*, chapter 7: Visual Modeling of Dynamic Gestures Using 3D Appearance and Motion Features, pages 103–120. Springer-Verlag, 2005.
- O1. **J. J. Corso**. Vision-Based Techniques for Dynamic, Collaborative Mixed-Realities. In B. J. Thompson, editor, *Research Papers of the Link Foundation Fellows*, volume 4. University of Rochester Press, 2004. Invited Report for Link Foundation Fellowship.

REPORTS (NOT PEER-REVIEWED) AND MANUSCRIPTS IN REVIEW

- R29. Y. Li, Y. Cao, and **J. J. Corso**. Instructional video generation. arXiv: 2412.04189, 2024.
- R28. F. Bellos, N. H. Nguyen, and **J. J. Corso**. VITRO: Vocabulary inversion for time-series representation optimization. arXiv: 2412.17921, 2024.
- R27. S. Storks, I. Bar-Yossef, Y. Li, Z. Zhang, **J. J. Corso**, and J. Chai. Explainable procedural mistake detection. arXiv: 2412.11927, 2024.

- R26. J. Marks, B. A. Griffin, and **J. J. Corso**. Class-wise autoencoders measure classification difficulty and detect label mistakes. arXiv: 2412.02596, 2024.
- R25. B. A. Griffin, J. Marks, and **J. J. Corso**. Zero-shot coreset selection: Efficient pruning for unlabeled data. arXiv: 2411.15349, 2024.
- R24. M. R. Ganesh, S. Y. Sekeh, and **J. J. Corso**. Q-TART: Quickly training for adversarial robustness and in-transferability. Technical Report 2204.07024, ARXIV, 2022.
- R23. A. Venkataraman, B. Griffin, and **J. J. Corso**. Learning kinematic descriptions using spare: Simulated and physical ARticultated extendable dataset. Technical Report 1803.11147, ARXIV, 2018.
- R22. V. Dhiman, S. Banerjee, B. Griffin, J. M. Siskind, and **J. J. Corso**. A critical investigation of deep reinforcement learning for navigation. Technical Report 1802.02274, ARXIV, 2018.
- R21. S. Patel, B. Griffin, K. Kusano, and **J. J. Corso**. Predicting future lane changes of other highway vehicles using rnn-based deep models. Technical Report 1801.04340, ARXIV, 2018.
- R20. M. R. Ganesh, E. Hofesmann, B. Min, N. Gafoor, and **J. J. Corso**. T-recs: Training for rate-invariant embeddings by controlling speed for action recognition. Technical Report 1803.08094, ARXIV, 2018.
- R19. E. Hofesmann, M. R. Ganesh, and **J. J. Corso**. M-PACT: An open source platform for repeatable activity classification research. Technical Report 1804.05879, ARXIV, 2018.
- R18. M. El Banani and **J. J. Corso**. Adviser networks: Learning what question to ask for human-in-the-loop viewpoint estimation. Technical Report 1802.01666, ARXIV, 2018.
- R17. R. Szeto, S. Stent, G. Ros, and **J. J. Corso**. A dataset to evaluate the representations learned by video prediction models. Technical report, ICLR Workshops, 2018.
- R16. S. Kumar, V. Dhiman, M. R. Ganesh, and **J. J. Corso**. Spatiotemporal articulated models for dynamic slam. Technical Report 1604.03526, arXiv, 2016.
- R15. **J. J. Corso**, A. Alahi, K. Grauman, G. D. Hager, L.-P. Morency, H. Sawhney, Y. Sheikh. Video Analysis for Body-worn Cameras in Law Enforcement. CCC Whitepaper. 2015.
- R14. A. Barbu, N. Siddharth, C. Xiong, **J. J. Corso**, C. D. Fellbaum, C. Hanson, S. J. Hanson, S. Hélie, E. Malaia, B. A. Pearlmutter, J. M. Siskind, T. M. Talavage, and R. B. Wilbur. The compositional natural of verb and argument representations in the human brain. Technical Report 1306.2293, arXiv, 2013.
- R13. C. S. Lea and **J. J. Corso**. Efficient hierarchical markov random fields for object detection on a mobile robot. Technical Report 1111.1599v1, arXiv, November 2011.
- R12. 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. J. Corso**, Y. Fu, and R. K. Srihari. GENIE TRECVID2011 multimedia event detection: Late-fusion approaches to combine multiple audio-visual features. In *NIST TRECVID Workshop*, 2011.
- R11. A. Y. C. Chen and **J. J. Corso**. Propagating multi-class pixel labels throughout video frames. In *Proceedings of Western New York Image Processing Workshop*, 2010.
- R10. J. A. Delmerico, **J. J. Corso**, and P. David. Boosting with stereo features for building facade detection on mobile platforms. In *Proceedings of Western New York Image Processing Workshop*, 2010.
- R9. H. Girgis and **J. J. Corso**. STP: The Sample-Train-Predict Algorithm and Its Application to Protein Structure Meta-Selection. Technical Report 2008-16, University at Buffalo SUNY, 2008.
- R8. I. Nwogu, **J. J. Corso**, and T. Bittner. The design of an ontology-enhanced anatomy labeler. Technical Report 2008-09, University at Buffalo SUNY, 2008.
- R7. C. Arnold, **J. J. Corso**, and A. Bui. An Unsupervised Approach to Automatic Image Annotation. In *NSF Biomedical Informatics Workshop: Expanding Secondary Use of Health Data*, 2007.
- R6. D. Burschka, G. Ye, **J. J. Corso**, and G. D. Hager. A Practical Approach for Integrating Vision-Based Methods into Interactive 2D/3D Applications. Technical report, The Johns Hopkins University, 2005. CIRL Lab Technical Report CIRL-TR-05-01.

- R5. **J. J. Corso**, M. Dewan, and G. D. Hager. Image Segmentation Through Energy Minimization Based Subspace Fusion. Technical Report CIRL-TR-04-01, The Johns Hopkins University, 2004.
- R4. **J. J. Corso**, N. Ramey, and G. D. Hager. Stereo-Based Direct Surface Tracking with Deformable Parametric Models. Technical report, The Johns Hopkins University, 2003. CIRL Lab Technical Report 2003-02.
- R3. **J. J. Corso**, G. Ye, D. Burschka, and G. D. Hager. Software Systems for Vision-Based Spatial Interaction. In *Proceedings of 2002 Workshop on Intelligent Human Augmentation and Virtual Environments*, pages D–26 and D–56, 2002.
- R2. **J. J. Corso** and G. D. Hager. Planar Surface Tracking Using Direct Stereo. Technical report, The Johns Hopkins University, 2002. CIRL Lab Technical Report.
- R1. **J. J. Corso** and J. D. Cohen. Out-Of-Core Voxelization of Large Scalar Fields for Interactive Multiresolution Volume Rendering. Technical report, The Johns Hopkins University, 2002. Graphics Lab Technical Report.

PATENTS

- P8. J. A. Marks, B. A. Griffin, B. E. Moore, **J. J. Corso**. Class-wise Autoencoders Measure Classification Difficulty and Detect Label Mistakes. Patent Pending (63/701,921) Submitted 10/2024.
- P7. B. A. Griffin, J. A. Marks, **J. J. Corso**. Blind Coreset Selection: Efficient Pruning for Unlabeled Data. Patent Pending (63/701,919) Submitted 10/2024.
- P6. K. D. Kusano, S. Patel, B. Griffin., **J. J. Corso**. System and Method for Vehicle Lane Change Prediction Using Structural Recurrent Neural Networks. No. US10611371B2.
- P5. J. Y. Song, S. Lemmer, **J. J. Corso**, W. S. Lasecki. Reconstructing 3D Video Using Particle Filtering to Aggregate Crowd Responses. No. US11361500B2.
- P4. D. Likosky, S. Yule, F. Pagani, M. Mathis, **J. J. Corso**, R. Daglius, E. M. Provost. Automated Identification and Grading of Intraoperative Quality. Patent Pending (16/705,371) Submitted 2019.
- P3. **J. J. Corso** and S. Sadanand. Method of Recognizing Activity in Video. Patent Pending (WO/2013/122675) Submitted 2013.
- P2. D. Das, M. Filipovich, **J. J. Corso**, G. D. Hager. Systems and Methods for Motion and Distance Measurement in Gastrointestinal Endoscopy. Patent Pending (13/457,305) Submitted 2013.
- P1. **J. J. Corso**, M. Smith, and M. Filipovich. System and Method for Mosaicing Endoscope Images Captured From Within A Cavity. Patent Pending (12/347,855) Submitted 2012.

ARTICLES IN MAINSTREAM MEDIA

- A2. **J. J. Corso**. *The open-source AI debate: Why selective transparency poses a serious risk*. VentureBeat. March 22, 2025.
<https://venturebeat.com/ai/the-open-source-ai-debate-why-selective-transparency-poses-a-serious-risk/>
- A1. **J. J. Corso** *Industry Voices — The Road to 99.999%: Eliminating Blind Spots in Automotive Visual AI*. Wards Auto. February 5, 2025
<https://www.wardsauto.com/software-defined-vehicles/industry-voices-the-road-to-99-999--eliminating-blind-spots-in-automotive-visual-ai>

ESSAYS

- E19. Understanding Dataset Difficulty with Class-Wise Autoencoders 2/2025
<https://medium.com/voxel51/understanding-dataset-difficulty-with-class-wise-autoencoders-dec813243e76>
- E18. Data Curation: The Beast Behind Every AI Model 10/2024
<https://medium.com/@jasoncorso/data-curation-the-beast-behind-every-ai-model-d136eac4da6b>
- E17. The Power of Open Source AI: How FiftyOne Drives the Future of Visual AI 10/2024
<https://medium.com/voxel51/the-power-of-open-source-ai-how-fiftyone-drives-the-future-of-visual-ai-74e7c83832ca>
- E16. So I Built This: Broadening the Impact of What You’ve Built in the Lab 6/2024
<https://medium.com/@jasoncorso/so-i-built-this-broadening-the-impact-of-what-youve-built->

in-the-lab-31a5e591713d

- E15. Observations in MLOps—A Fragmented Mosaic of Mismatched Expectations 4/2024
<https://medium.com/@jasoncorso/observations-on-mlops-a-fragmented-mosaic-of-mismatched-expectations-3488685ec0b6>
- E14. Getting Rid of the Bad Seeds: A Quick Intro to Seed Rejection in Human+AI Teams 3/2024
<https://medium.com/@jasoncorso/getting-rid-of-the-bad-seeds-a-quick-intro-to-seed-rejection-in-human-ai-teams-7497129323ff>
- E13. From the Kitchen to the Lab: Why Cooking Became AI's Favorite Dish 3/2024
<https://medium.com/@jasoncorso/from-the-kitchen-to-the-lab-why-cooking-became-ais-favorite-dish-2c089017c12a>
- E12. How to Effectively Mentor New Technical Writers, Part 2: Selection 3/2024
<https://medium.com/@jasoncorso/how-to-effectively-mentor-new-technical-writers-part-2-selection-5d9c89f23a55>
- E11. How to Effectively Mentor New Technical Writers, Part 1: The Writing 3/2024
<https://medium.com/@jasoncorso/how-to-effectively-mentor-new-technical-writers-part-1-the-writing-fe0ffbc17e1e>
- E10. A Window into Virtual Open Office Hours and Its Unseen Benefits 2/2024
<https://medium.com/@jasoncorso/a-window-into-virtual-open-office-hours-and-its-unseen-benefits-4d7f163a1e45>
- E9. How Developer Tooling Continues to Shape the AI Landscape 2/2024
<https://medium.com/@jasoncorso/how-developer-tooling-continues-to-shape-the-ai-landscape-c6a71142aa3d>
- E8. How to Negotiate Your First Faculty Job Offer 2/2024
<https://medium.com/@jasoncorso/how-to-negotiate-your-first-faculty-job-offer-852beed1dd73>
- E7. Is Open Source AI Bull? 2/2024
<https://medium.com/@jasoncorso/is-open-source-ai-bull-9da010411658>
- E6. How to Read Conference Papers 1/2024
<https://medium.com/@jasoncorso/how-to-read-conference-papers-fa78c75f78aa>
- E5. Navigating Academia to Startup: What Kind of Companies Do Scholars Create? 1/2024
<https://medium.com/@jasoncorso/navigating-academia-to-startup-what-kind-of-companies-do-scholars-create-9fe5b7e44fbc>
- E4. Annotation is Dead 1/2024
<https://medium.com/@jasoncorso/annotation-is-dead-1e37259f1714>
- E3. The Machine Learning Random Walk 10/2023
<https://medium.com/@jasoncorso/the-machine-learning-random-walk-0739a38bdc54>
- E2. How to Get the most out of CVPR 6/2023
<https://medium.com/voxel51/how-to-get-the-most-out-of-cvpr-4abaa09dbc50>
- E1. Too Many Pixels, So Little Time: Why You Need Data-Centric Tooling in Your AI Stack 6/2023
<https://medium.com/voxel51/too-many-pixels-so-little-time-why-you-need-data-centric-tooling-in-your-ai-stack-fb888ea704fa>

- PODCAST INTERVIEWS
- I13. Cooperative Systems, Data Transparency & Quality, and the Year of Small AI 4/2025
on the *Open Source Data Podcast* hosted by Charna Parkey. S7E5.
<https://www.datastax.com/resources/podcast/cooperative-systems-data-transparency-and-quality-and-the-year-of-small-ai-or-dr-jason-corso>
<https://www.youtube.com/watch?v=zWRXQr08uK8>
 - I12. Achieving 99.999% Accuracy for Visual AI 2/2025
on the *Techstrong.TV Digital CXO Podcast* hosted by Amanda Razani. E110.
<https://techstrong.tv/videos/digital-cxo-podcast/achieving-99-999-accuracy-for-visual-ai-digital-cxo-podcast-ep110>
 - I11. Computer Vision in Action 12/2024
on *UX Magazine's Invisible Machines* hosted by Robb Wilson and Josh Tyson. S4E15
<https://uxmag.com/podcast/computer-vision-in-action>

- I10. AI Partnerships Drive Developer Innovation co-interviewed with Paula Ramos (Intel) 10/2024
on the *Insight.Tech Podcast* hosted by Christine Cardoza
<https://brnw.ch/21wnrlg>
- I9. Image Processing and Robotics 8/2024
on the *Neural Network Podcast* hosted by A Priyadarshan and V Lakshmanan
<https://open.spotify.com/episode/5y8A6NYxUXVY5wmTmsN17L?si=Pf6UKqYdRueYKSAszav55w>
- I8. Professor-Student Duo Takes on AI 5/2024
on the *How I Met My Co-Founder: The Podcast* hosted by Annie Garofalo. E1.8
<https://creators.spotify.com/pod/profile/how-i-met-my-cofounder/episodes/ep-1-8--professor-student-duo-takes-on-ai--interview-with-voxel51-e2j4i05/a-ab7nqm6>
- I7. \$15.5 Million Raised to Build the Future of Developer Tools for MLOps 4/2024
on *Frontlines: Category Visionaries* hosted by Brett Strapper. E442.
<https://www.frontlines.io/podcasts/jason-corso/>
- I6. AI Data Accuracy—How to Approach Working with Datasets 4/2024
on *Chaos to Clarity* hosted by Eric Weiss.
<https://open.spotify.com/episode/0jvj0jjqepi96gjlomonam>
- I5. Shaping the Future of AI: From Academia to Entrepreneurship with Jason Corso 2/2024
on *Humans of AI* hosted by Sheikh Shuvo. E39.
<https://humansofai.transistor.fm/episodes/39-shaping-the-future-of-ai-from-academia-to-entrepreneurship-with-jason-corso>
- I4. From Academic Roots to Open Source to Tech Entrepreneurship 1/2024
on *Skilltorch Disrupt* hosted by Milica Mishiri.
<https://open.spotify.com/episode/3buivfs2jichlcwyy62224>
- I3. ML Pulse Report with Voxel51 CSO Jason Corso and SAMA VP Duncan Curtis 11/2023
on *How AI Happens* hosted by Rob Stevenson. E84.
<https://podcasts.apple.com/us/podcast/ml-pulse-report-with-voxel51-cso-jason-corso-and-sama/id1568452184?i=1000635708980>
- I2. The Academic Side of AI 10/2023
on *The Peggy Smedley Show* hosted by Peggy Smedley. E841.
<https://peggysmedleyshow.com/the-academic-side-of-ai>
- I1. Research and Tooling for Computer Vision Systems with Jason Corso 1/2021
on *Datacast* hosted by James Le. E51
<https://jameskle.com/writes/jason-corso>

QUOTED IN MEDIA

- Q14. AI-powered self-service: Ethical, privacy considerations 3/25/2025
in *Retail Customer Experience*, by Sandra Carpenter
<https://www.retailcustomerexperience.com/articles/ai-powered-self-service-ethical-privacy-considerations/>
- Q13. Google, OpenAI target state laws in AI action plan 3/21/2025
in *TechTarget*, by Makenzie Holland
<https://www.techtarget.com/searchenterpriseai/news/366621060/Google-OpenAI-target-state-laws-in-AI-action-plan>
- Q12. Tech leaders sound alarm over DOGE's AI firings, impact on American talent pipeline, in *ZDNet*, by Radhika Rajkumar 3/17/2025
<https://www.zdnet.com/article/tech-leaders-sound-alarm-over-doges-ai-firings-impact-on-american-talent-pipeline/>
- Q11. Next gen kiosks innovate across industry applications 2/28/2025
in *Kiosk Marketplace*, by Sandra Carpenter
<https://www.kioskmarketplace.com/articles/beyond-the-touchscreen-how-next-gen-kiosks-are-reshaping-our-world/>
- Q10. Potential cuts at AI Safety Institute stoke concerns in tech industry 2/25/2025
in *The Hill*, by Miranda Nazzaro
<https://thehill.com/policy/technology/5161687-potential-cuts-at-ai-safety-institute-stoke-concerns-in-tech-industry/>
- Q9. Is DeepSeek next in line for a TikTok-like U.S. ban? 2/12/2025
in *NIKKEI Asia*, by Yifan Yu
<https://asia.nikkei.com/Business/Technology/Tech-Asia/Is-DeepSeek-next-in-line-for-a->

TikTok-like-U.S.-ban

- Q8. How a Free Software Strategy Catapulted DeepSeek to AI Stardom 1/29/2025
in *The Wall Street Journal*, by Belle Lin
<https://www.wsj.com/articles/how-a-free-software-strategy-catapulted-deepseek-to-ai-stardom-67e98349>
- Q7. AI takes off in Washtenaw County in *Concentrate*, by Sarah Rigg 1/8/2025
<https://www.secondwavemedia.com/concentrate/features/ai-takes-off-in-washtenaw-county.aspx>
- Q6. Enterprise AI predictions 2025: The year GenAI creates tangible value? 12/10/2024
in *Verdict*, by Lara Williams
<https://www.verdict.co.uk/enterprise-ai-predictions-2025-the-year-genai-creates-tangible-value/>
- Q5. Automotive AI: Applications that come with challenges in creation. 12/9/2024
in *Just Auto*, by Frankie Youd
<https://www.just-auto.com/interview/automotive-ai-applications-that-come-with-challenges-in-creation/>
- Q4. What will Trump's new term mean for A.I.? in *Newsweek*, by Marie Boran 11/17/2024
<https://www.newsweek.com/what-will-trumps-new-term-mean-ai-1982615>
- Q3. Trump revoking Biden AI EO will make industry more chaotic, experts say 11/14/2024
in *VentureBeat*, by Emilia David
<https://venturebeat.com/ai/trump-revoking-biden-ai-eo-will-make-industry-more-chaotic-experts-say/>
- Q2. What to expect for AI under Trump in *The Hill*, by Miranda Nazzaro 11/12/2024
<https://thehill.com/policy/technology/4986557-trump-ai-industry-impact/>
- Q1. Newsom throws AI regulation fight into uncertainty with veto 10/2/2024
in *The Hill*, by Miranda Nazzaro
<https://thehill.com/policy/technology/4910138-newsom-throws-ai-regulation-fight-into-uncertainty-with-veto/>

SOFTWARE &
DATA SETS

- FiftyOne** 2020–2025
Open-Source Tool for CV/ML dataset visualization, analysis and model debugging.
<https://github.com/voxel51/fiftyone>
Reached 3,000,000 installs as of 12/2024.
- Unified VLP** 2020
Code to accompany our AAAI 2020 paper on unified video-language pretraining.
<https://github.com/LuweiZhou/VLP>
- Player51** 2019
Open-Source Tool for efficient, web-based video rendering with CV/ML labels.
<https://github.com/voxel51/player51>
- BubbleNets** 2019
Code to accompany our CVPR 2019 paper on intelligence selection of the frame to annotate in video object segmentation.
<https://github.com/griffbr/BubbleNets>
- ETA** 2018
Open-Source extensible toolkit for video analytics.
<https://github.com/voxel51/eta>
- ViP** 2019
Video software platform for object detection and action recognition based on PyTorch.
<https://arxiv.org/abs/1910.02793>
<https://github.com/MichiganC0G/ViP>
- Grounded Video Description** 2019
Code and Dataset to accompany our CVPR 2019 oral paper on grounded video description.
Code: <https://github.com/facebookresearch/grounded-video-description>
Data: <https://github.com/facebookresearch/ActivityNet-Entities>
- TASED-Net** 2019
Code to accompany our ICCV 2019 paper on temporally-aggregating spatial encoder-decoder network for video saliency detection.

<https://github.com/MichiganCOG/TASED-Net>

M-PACT 2018
General purpose software framework for all video processing and activity recognition based on TensorFlow.

<https://arxiv.org/abs/1804.05879>

<https://github.com/MichiganCOG/M-PACT>

YouCook2 2017
Large-scale data set of third-person cooking videos categorized 89 recipes with each recipes having at least 21 unique instructional videos. Time-coded human-written recipes are provided along with the raw video.

<http://youcook2.eecs.umich.edu/>

Click-Here CNNs 2017
Code to accompany our ICCV 2017 paper on human-in-the-loop inference for keypoint-based guidance of monocular viewpoint estimation.

<https://github.com/rszeto/click-here-cnn>

End-To-End Video Captioning 2016
Our system for end-to-end video captioning using text-conditional semantic attention.

<https://github.com/LuweiZhou/e2e-gLSTM-sc>

Video2Text.net 2013
A website and web-service for automatic conversion of videos to natural language sentences based on the video content. This website showcases our work in the vision+language domain.

<http://www.video2text.net>

YouCook 2013
Data set of third-person cooking videos categorized into six styles of cooking and selected from open-source web videos of different kitchens and complexity levels. It contains object and action bounding boxes as well as multiple natural language descriptions of each video.

<http://web.eecs.umich.edu/~jjcorso/r/youcook>

Random Forest Distance 2012
Software implementation of our KDD 2012 tree-structured metric learning paper.

http://web.eecs.umich.edu/~jjcorso/pubs/RFD_Package.zip

LIBSVX 2012
Supervoxel library: a set of methods for early video processing by computing supervoxel segmentations as well as a quantitative benchmark for fair comparisons of those segmentations.

<http://web.eecs.umich.edu/~jjcorso/r/supervoxels>

Winner Best Demo Prize at 2nd Greater New York Multimedia and Vision Meeting.

6/2012

Winner Best Open Source Code 3rd Prize at IEEE CVPR 2012.

6/2012

Action Bank 2012
Code to compute a high-level feature representation of activity in video.

<http://web.eecs.umich.edu/~jjcorso/r/actionbank>

Winner Best Open Source Code 3rd Prize at IEEE CVPR 2012.

6/2012

Chen Xiph.org 2011
Data set to support video label propagation and video semantic segmentation. Contains 8 videos densely labeled (frame-by-frame) into the 21 MSRC classes. The videos are selected from the xiph.org repository.

<http://www.cse.buffalo.edu/~aychen/LabelPropagation/labelpropagation.zip>

Video Label Propagation 2011
Code to propagate an initial segmentation through a video sequence.

<http://www.cse.buffalo.edu/~aychen/LabelPropagation/propagatelabel.m>

MuleSeg 2006
Extensible software for multilevel segmentation of 2D and 3D images based on an extended Segmentation by Weighted Aggregation algorithm using the Bayesian model-aware affinity

GUSTO 2002
System for interactive, hierarchical rendering of large (out-of-core) 3D scalar fields, including unstructured grids, structured grids, and voxels. Jointly with Joshua Leven, Jonathan D. Cohen, and Subodh Kumar.

XVision2 2001
Modular software architecture for real-time vision development. Jointly with Gregory Hager, Darius Burschka, Sam Lang, and Xiangtian Dai.

FUNDING

TOTAL ACADEMIC FUNDING: \$55,731,236 (\$43,132,075 as PI)

SHARE CREDIT: \$19,947,659

Funding is sorted by start date (recent first).

Active funding (36, 40, and 41) is prefixed with an asterisk.

Dollar amounts listed as \$Total (\$Share; Percent Credit)

- F41. ***PI: VIGIL—Vectors of Intelligent Guidance** \$24,595,343 (\$5,558,548; 23%)
in Long-Reach Rural Healthcare
 SOURCE: ARPA-H PARADIGM 11/2024–11/2029
 COLLABORATORS: Alison Arnold (Central Mich RHEI), Nathaniel Blanchard (Colorado State CS), Bruce Draper (Colorado State CS), Enrique Dunn (Stevens CS), Ehsan Elhamifar (North-eastern CS), Anhong Guo (UMich CSE), Nikhil Krishnaswamy (Colorado State CS), Donald Likosky (UMich Cardiac Surgery), Emily J Mackay (UPenn Anesthesiology and Critical Care), Prashant Mahajan (UMich Emergency Medicine and Pediatrics), Milisa Manojlovich (UMich Nursing), Muntaha Samad (RTX Raytheon BBN), Andrew Owens (UMich ECE), Francis Pagani (UMich Cardiac Surgery), Alison Marie Pouch (UPenn Radiology and Bioengineering), Sethu K Reddy (Central Mich Endocrinology), Jeff Siskind (Purdue ECE), Sarath Sreedharan (Colorado State CS), Steve Vance (Central Mich Clinical Education), Brain van Voorst (RTX Raytheon BBN), Kevin Ward (UMich Emergency Medicine and Biomedical Engineering), Chenliang Xu (Rochester CS)
 OBJECTIVE: Investigate AI task guidance to enable the upskilling of medical professionals in rural healthcare settings.
- F40. ***Co-I: QUASAR—Quantified Uncertainty Anchors for Same-Day** \$5,913,977 (\$2,219,899; 37.5%)
Autonomy and Rapid Transfer through Dense Learning
 SOURCE: DARPA TIAMAT 7/2024–7/2027
 COLLABORATORS: Henry Liu (PI; UMich CEE), Maani Ghaffari (UMich NAME), Ram Vasudevan (UMich ROB)
 OBJECTIVE: Investigate computational learning models that can leverage semantic information in scene to support more capable transfer from low to high fidelity situations in simulation and real-world scenarios.
- F39. **PI: Passive Ranging in Low-Light Conditions using Fused EO/IR** \$59,100 (\$59,100; 100%)
Multi-View Imaging
 SOURCE: ARO STTR (Sub from ThirdInsight; Phase 1) 4/2024–12/2024
 OBJECTIVE: Investigate the fusion of optical and infrared imagery to estimate depth of objects in view.
- F38. **PI: MILLY—Multi-directional Loosely-Linked Archetype Models** \$5,913,977 (\$2,219,899; 37.5%)
for Perceptually-enabled Task Guidance
 SOURCE: DARPA PTG 1/2022–1/2025
 COLLABORATORS: Joyce Chai (UMich CSE), Jeff Siskind (Purdue ECE), Chenliang Xu (Rochester CS), Enrique Dunn (Stevens CS)
 OBJECTIVE: Investigate joint learning of knowledge and attention, grounded in real-world perception for agent-based guidance of purposeful tasks.
- F37. **Co-I: Low-Cost 3D Perception for Mobile Manipulation** \$1,851,143 (\$108,678; 5.8%)
in Unstructured Human Environments
 SOURCE: TRI 4/2021–3/2024
 COLLABORATORS: Brent Griffin (UMich ECE, PI), Justin Johnson (UMich CSE), David Fouhey (UMich CSE)
 OBJECTIVE: Investigate new methods for adaptive 3D perception in support of complex mobile manipulation tasks on home service robotic platforms.
- F36. ***Co-I: Novel Assessments of Technical and Non-Technical** \$3,443,788 (\$391,465; 11%)
Cardiac Surgery Quality
 SOURCE: NIH 12/2019–11/2025
 COLLABORATORS: Donald Likosky (MED, PI), Frank Pagani (MED, PI), Matthew Caldwell (MED), Sarah Krein (MED), Milisa Manojlovich (Nursing), Michael Mathis (MED), Min Zhang (Biostats), Steven Yule (BWH), Roger Dias (BWH)
 OBJECTIVE: Advance the modeling and understanding of technical and non-technical characteristics that correlate with various surgical outcomes.
- F35. **Co-I: Interactive Learning for Manipulating Piles of Stuff** \$1,193,108 (\$270,357; 23%)
 SOURCE: TRI 7/2018–12/2020
 COLLABORATORS: Brent Griffin (ECE, PI), Dmitry Berenson (ECE)
 OBJECTIVE: Investigate new methods for the intersection of computer vision and mobile

manipulation, with an emphasis on interactively learning perceptual models.

- F34. **Co-PI:** *Human-Augmented Computer Vision for Robust 3D Scene Reconstruction at Scale* \$504,045 (\$257,394; 51%)
SOURCE: TRI 1/2018–12/2020
COLLABORATORS: Walter Lasecki (CSE)
OBJECTIVE: Understand how humans can be involved in the loop at inference time in a hybrid intelligence mindset to reconstruct complicated visual scenes from monocular data.
- F33. **PI:** *TARDIS-V: Multi-Layer Compositional Activity Understanding using Joint Visual and Semantic Inference* \$1,439,890 (\$949,667; 66%)
SOURCE: IARPA DIVA (sub to SRI) 10/2017–5/2021
COLLABORATORS: Jia Deng (CSE)
OBJECTIVE: Investigate new methods for fine-grained human activity recognition in ground-level and first-person-view cameras.
- F32. **PI:** *BOCA: Body-Worn Camera Analytics* \$681,290 (\$681,290; 100%)
SOURCE: NIST PSIAP 6/2017–5/2019
COLLABORATORS: Yan Tom Yan (Texas State University)
OBJECTIVE: Investigate new methods for transfer learning between third-person and first-person perspective human activity video.
- F31. **PI:** *SPIDER: Subspace Primitives that are Interpretable and Diverse* \$1,596,900 (\$798,450; 50%)
SOURCE: DARPA D3M 4/2017–3/2021
COLLABORATORS: Laura Balzano (ECE)
OBJECTIVE: Investigate new machine learning models and algorithms based on subspace primitive ideas and involving multimodal data sources and explicit invariance properties.
- F30. **PI:** *CI:New: COVE—Computer Vision Exchange for Data, Annotation, and Tools* \$343,000 (\$343,000; 100%)
SOURCE: National Science Foundation (NSF) CRI 8/2016–7/2019
COLLABORATORS: Kate Saenko (Boston University) and Walter Scheirer (Notre Dame)
OBJECTIVE: Develop a website that makes data and related resources openly and widely available to the computer vision research community.
- F29. **Co-PI:** *Efficient Human-in-the-Loop Computer Vision Algorithms to Create Datasets of Rare Traffic Events From Video* \$199,810 (\$89,915; 50%)
SOURCE: UM Mobility Transformation Center 1/2017–12/2017
COLLABORATORS: Walter Lasecki (CSE)
OBJECTIVE: Investigate improved strategies for incorporating humans in the loop for monocular reconstruction problems.
- F28. **Co-PI:** *Task-Oriented Active Perception and Motion Planning for Manipulating Piles of Stuff* \$499,184 (\$249,592; 50%)
SOURCE: Toyota Research Institute (TRI) 1/2017–6/2018
COLLABORATORS: Dmitry Berenson (ECE), Brent Griffin (ECE)
OBJECTIVE: Investigate tightly coupled active perception and planning strategies for interactive manipulation and task-oriented processing of assortments of objects in tabletop settings.
- F27. **PI:** *MEDISPHERE* \$650,278 (\$650,278; 100%)
SOURCE: DARPA MediFor (Sub from PAR Government) 9/2016–8/2018
COLLABORATORS: Brent Griffin (ECE)
OBJECTIVE: Support the creation of manipulated video data for the forensics TA1 teams on the program.
- F26. **PI:** *Computer Vision and Crowdsourcing for Vehicle Crash Analysis* \$138,324 (\$69,162; 50%)
SOURCE: Denso 3/2016–8/2016
COLLABORATORS: Walter Lasecki (CSE)
OBJECTIVE: Develop a system where human can operate in a monocular reconstruction loop with an emphasis on road scenes.
- F25. **PI:** *RobotSLANG: Simultaneous Localization, Mapping, and Language Acquisition* \$649,999 (\$649,999; 100%)
SOURCE: National Science Foundation (NSF) NRI 9/2015–8/2020
COLLABORATORS: Jeffrey Mark Siskind (Purdue ECE, Separate NSF Collaborative Grant)
OBJECTIVE: Generalize simultaneous localization and mapping to include language acquisition for unknown environments and unknown symbol grounding functions.

- F24. **Co-PI:** *Novel trans-synaptic tracers and computational tools for labeling and automated reconstructing dense neural circuits at single cell resolution in the mouse brain* \$100,000 (\$50,000; 50%)
 SOURCE: University of Michigan MiBrain Initiative 8/2015-7/2016
 OBJECTIVE: Integrate computer vision and machine learning tools to label, reconstruct and classify all the neurons that connect to a single subtype-defined neuron at single cell and single synapse resolution in the mouse brain.
- F23. **PI:** *Blind Multibody Prediction with Semantic Dynamical Models* \$336,524 (\$488,060; 100%)
 SOURCE: Toyota 12/2015-2/2018
 OBJECTIVE: Model multibody dynamics to predict the safe egress time window for sensor-blind autonomous or assistive vehicles.
- F22. **PI:** *Learning from Unconstrained Instructional Videos with Structured and Sparse Models* \$100,000 (\$50,000; 50%)
 SOURCE: Samsung GRO 10/2015-9/2016
 COLLABORATORS: Cornelia Fermuller (U Maryland, College Park, CS)
 OBJECTIVE: To develop machine learning methods that can automatically learn the semantic human-object relationships of temporal processes from unstructured web-videos.
- F21. **PI:** *Action Co-Discovery As A Cross-Reconstruction Problem* \$290,481 (\$290,481; 100%)
 SOURCE: Army Research Office (ARO) 8/2015-7/2019
 OBJECTIVE: Formulate, solve and study the detection of common actions across multiple unsupervised videos as a cross-reconstruction problem.
- F20. **PI:** *Vehicle Tethered Quadrotor Control* \$25,000 (\$12,500; 50%)
 SOURCE: TARDEC 3/2015-9/2015
 COLLABORATORS: Necmiye Ozay (ECE) and Dimitra Panagou (AERO)
 OBJECTIVE: Determine the feasibility of GPS-denied control methods for tethered quadrotor operation from a vehicle.
- F19. **PI:** *Pictorial Structures for Scene Context* \$620,000 (\$620,000; 100%)
 SOURCE: DARPA STTR (Sub from Soartech; Phases I & II) 3/2015-9/2019
 OBJECTIVE: Integrate the SOAR cognitive architecture and state of the art computer vision for improving robustness and efficiency in scene understanding.
- F18. **PI:** *Computational PB&J: Learning from Instructional Video* \$67,657 (\$67,657; 100%)
 SOURCE: Google Faculty Research Program 4/2015-4/2018
 OBJECTIVE: This project investigates the rich problem of learning visual- and language-grounded process models from instructional video content.
- F17. **PI:** *CI-New: Collaborative Research: Federated Data Set Infrastructure for Recognition Problems in Computer Vision* \$150,000 (\$75,000; 50%)
 SOURCE: NSF Community Research Infrastructure 10/2014-10/2017
 COLLABORATORS: Kate Saenko (UMass Lowell, CS)
 OBJECTIVE: This project seeks to establish a federated infrastructure for datasets and annotations within the computer vision community
- F16. **PI:** *Transferring ACE to the Analyst* \$250,000 (\$250,000; 100%)
 SOURCE: DARPA CSSG Phase III 7/2012-7/2014
 OBJECTIVE: The main objective is to investigate how our developments in the Phase II active clustering into real application use with our defense partner.
- F15. **PI:** *Computer Vision and Mobile Robot Technologies for Advanced Emergency Response* \$261,178 (\$261,178; 100%)
 SOURCE: CUBRC (Federal Highway Administration) 10/2011-3/2014
 OBJECTIVE: This project seeks to develop and deploy single and multiple robotic platforms capable of concurrently mapping emergency sites in normal and adverse conditions.
- F14. **Co-PI:** *Objective Imaging-Based Assessment of Smoking Behavior from Used Filters* \$275,000 (\$164,313; 60%)
 SOURCE: National Institutes of Health NCI 1 R21 CA160825-01 9/2011-9/2014
 COLLABORATORS: O'Connor (PI, Roswell Park)
 OBJECTIVE: The proposal seeks to refine our ability to quantify the smoking behavior through digital image analysis of cigarette filters.
- F13. **PI:** *Two-Rank Mobile Robot Fleet for Swarm Surveillance, Warfighter Assistance, and other Army-related Research and Research-Related Education* \$250,000 (\$200,000; 80%)
 SOURCE: Army Research Office DURIP 9/2011-6/2013

COLLABORATORS: Demirbas (CSE), Fu (CSE), Krovi (MAE), Nagi (MAE)
 OBJECTIVE: The basic goal is to enhance the DoD's capabilities for using fleets and swarms of ground-based sensor-rich unmanned and autonomous mobile robots (UGVs).

- F12. **PI:** *GBS: Guidance By Semantics—Using High-Level Visual Inference to Improve Vision-based Mobile Robot Localization* \$150,000 (\$150,000; 100%)
 SOURCE: Army Research Office Young Investigator Program 6/2011-6/2014
 OBJECTIVE: The goal is to investigate how semantic perception can be used to improve the accuracy, speed, and robustness of vision-based localization of mobile robot platforms.
- F11. **PI:** *Comprehensive Object Detection Library for Large-Scale Image Analytics* \$190,975 (\$190,975; 100%)
 SOURCE: Naval Postgraduate School 5/2011-10/2012
 OBJECTIVE: The goal of the proposed work is to build a comprehensive object detection library and evaluate the methods for large-scale image sets.
- F10. **PI:** *Ontology, Event Agents and Event Recounting for ALADDIN* \$999,469 (\$799,575; 80%)
 SOURCE: IARPA (sub from Kitware, Inc.) 3/2011-3/2016
 COLLABORATORS: Fu (CSE) and R. Srihari (CSE)
 OBJECTIVE: The goal of this project is to improve the representation and indexing of objects and events in large-scale video analysis by efficiently encoding the low-level perceptual entities in the video and grounding them with rich high-level semantics.
- F9. **PI:** *ISTARE: Intelligent Spatio-Temporal Activity Reasoning Engine* \$2,208,368 (\$1,104,184; 50%)
 SOURCE: DARPA Mind's Eye Program 9/2010-12/2013
 COLLABORATORS: Fu (CSE), Ceusters (Psychiatry), Krovi (MAE), Petropoulos (CSE)
 OBJECTIVE: The goal of this proposal is a development of a methodology for representation, learning, recognition of and reasoning over activities in persistent surveillance videos.
- F8. **Co-PI:** *II-NEW: Acquisition of a Biomedical Computing Infrastructure* \$588,554 (\$58,855; 10%)
 SOURCE: National Science Foundation CRI 9/2010-9/2011
 COLLABORATORS: Chaudhary (PI, CSE), Hoffmann (NS), Krovi (MAE), Furlani (CCR)
 OBJECTIVE: The major goal of this project is to establish a state-of-the-art computing infrastructure for memory- and compute-bound problems in biomedicine.
- F7. **PI:** *Semantic Video Summarization With Ontology-Driven Probabilistic Inference on Massive Multimedia Collections* \$357,080 (\$357,080; 100%)
 SOURCE: CIA/IC Postdoc Fellowship Program 7/2010-7/2013
 OBJECTIVE: The main goal of the proposal is to advance the understanding of how probabilistic ontologies of semantic entities and entity-entity relationships can drive inference for semantic summarization of content in massive video collections.
- F6. **PI:** *ACE – Active Clustering for Exploitation and Defense Forensics* \$399,780 (\$399,780; 100%)
 SOURCE: DARPA CSSG Phase II 6/2010-6/2012
 OBJECTIVE: The main objective is to investigate how active clustering can be used to induce high-level models of phenomena in video with the help of a user.
- F5. **PI:** *Digital Imaging of Cigarette Filters* \$27,958 (\$27,958; 100%)
 SOURCE: Health Research, Inc. 9/2009-12/2009
 OBJECTIVE: The major goal of this project is to develop improved image analysis techniques for quantifying the properties of cigarette filters in digital images.
- F4. **PI:** *CAREER: Generalized Image Understanding with Probabilistic Ontologies and Dynamic Adaptive Graph Hierarchies* \$539,086 (\$539,086; 100%)
 SOURCE: National Science Foundation CAREER Program 7/2009-7/2014
 OBJECTIVE: The major research goal is to investigate a unified model of image representation integrating probabilistic methods, machine learning, probabilistic ontologies, and dynamic adaptive graphs to advance our ability to solve the image understanding problem.
- F3. **PI:** *Probabilistic Ontology Induction for Generalized Video Understanding* \$99,670 (\$99,670; 100%)
 SOURCE: DARPA CSSG Phase I 4/2009-7/2010
 OBJECTIVE: The major goal is to develop the foundational models for learning high-level semantic representations of objects from video data, automatically and unsupervisedly.
- F2. **Co-PI:** *Multimodal Command-and-Control By Integrating Two-Handed Gestures and Speech* \$129,953 (\$32,488; 25%)
 SOURCE: Hewlett Packard Labs Innovation Research Program 8/2008-8/2010

COLLABORATORS: Govindaraju (PI)

OBJECTIVE: The major goal of this project is to integrate multimodal processing into robust gesture recognition.

- F1. **Co-PI: Reconstructing CT Images from a Limited Number of Projections** \$27,000 (\$9,000; 33%)
SOURCE: UB Interdisciplinary Research Development Fund
COLLABORATORS: Xu (PI, CSE) and Hoffman (NS)
OBJECTIVE: This project explores how redundant information in different projections can be combined with a priori knowledge of the anatomy to reduce dosage while maintaining high quality reconstructions.

AFFILIATIONS	Senior Member , Institute of Electrical and Electronics Engineers (IEEE) Member since 1998	2016-PRESENT
	Member , Association for the Advancement of Artificial Intelligence (AAAI)	2011-PRESENT
	Member , Mathematical Association of America (MAA)	2005-PRESENT
	Member , IEEE Robotics and Automation Society	2003-2005
	Member , Upsilon Pi Epsilon, Computer Science Honor Society	1999-PRESENT
	Member , Alpha Sigma Nu, National Jesuit Honor Society	1999-PRESENT
	Member , Association of Computing Machinery (ACM)	1998-PRESENT
	Member , IEEE Computer Society (IEEE)	1998-PRESENT

SERVICE

EDITORSHIPS / EDITORIAL BOARD MEMBERSHIP

- * **Associate Editor**, IEEE Trans. on Pattern Analysis and Machine Intelligence 10/2014–12/2019
Published by IEEE
- * **Editorial Board Member**, International Journal of Computer Vision 8/2014–3/2019
Published by Springer
- * **Associate Editor**, Computer Methods and Programs in Biomedicine 9/2009–6/2014
Published by Elsevier

CONFERENCE / WORKSHOP CHAIR

- * **General Chair**, FIVER 2018: Fine-Grained Instructional Video Understanding Workshop (at CVPR 2018) 2018
Jointly with Ivan Laptev, Josef Sivic and Luowei Zhou
- * **General Chair**, Midwest Computer Vision Workshop 2018
Jointly with Jia Deng
- * **General Chair**, BigVision 2016: International Workshop on Large Scale Visual Recognition and Retrieval (at CVPR 2016) 2016
Jointly with Samy Bengio (Google) and Fei-Fei Li (Stanford)
- * **General Chair**, The Future of Datasets in Vision 2015 (at CVPR 2015) 2015
Jointly with Kate Saenko

SENIOR PROGRAM COMMITTEE

- * **Area Chair**, IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) 2025
- * **Area Chair**, IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) 2024
- * **Associate Editor**, IEEE Intl. Conf. on Robotics and Automation (ICRA) 2024
- * **Area Chair**, IEEE Intl. Conf. on Computer Vision 2023
- * **Area Chair**, IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) 2021
- * **Area Chair**, European Conference on Computer Vision (ECCV) 2020
- * **Area Chair**, IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) 2019
- * **Area Chair**, IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) 2017
- * **Associate Editor**, IEEE Intl. Conf. on Robotics and Automation (ICRA) 2017
- * **Area Chair**, AAAI Conference on Artificial Intelligence 2017
- * **Area Chair**, AAAI Conference on Artificial Intelligence 2016
- * **Associate Editor**, IEEE Intl. Conf. on Robotics and Automation (ICRA) 2016
- * **Program Chair**, BigVision 2015: International Workshop on Large Scale Visual Recognition and Retrieval (at CVPR 2015) 2015
Jointly with O. Russakovsky (Stanford), J. Deng (Michigan) and Y. Lin (NEC)
- * **Associate Editor**, IEEE Intl. Conf. on Robotics and Automation (ICRA) 2015
- * **Area Chair**, IEEE Intl. Conf. on Computer Vision 2015
- * **Program Chair**, BigVision 2014: Intl. Workshop on Large Scale Visual Recognition and Retrieval (at CVPR 2014) 2014

- Jointly with J. Deng (Michigan), A. Berg (UNC) and Y. Lin (NEC)*
- * **Area Chair**, European Conference on Computer Vision (ECCV) 2014
 - * **Area Chair**, IEEE Winter Conf. on Applications of Computer Vision (WACV) 2014
 - * **Area Chair**, IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) 2013
 - * **Area Chair**, IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) 2012

OTHER CONFERENCE AND WORKSHOP SERVICE

- * **Workshops Chair**, IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017
- * **Video Chair**, IEEE International Conference on Computer Vision (ICCV) 2015
- * **Organizing Committee**, CONTACT 2014: First International Workshop on Computer vision + ONTology Applied Cross-disciplinary Technologies (at ECCV 2014) 2014
Jointly with M. Cristani (Università degli Studi di Verona) and R. Ferrario (Italian National Research Council)
- * **Student Activities Chair**, IEEE Conf. on Computer Vision and Pattern Recognition 2014
- * **Organizing Committee**, ACCV Workshop on Detection and Tracking in Challenging Environments (DTCE 2012) 2012
Jointly with J. Lim (Hanyang Univ.), B. Han (POSTECH), B. Triggs (CNRS) and A. Elgammal (Rutgers).
- * **Student Activities Chair**, IEEE Conf. on Computer Vision and Pattern Recognition 2012
- * **Organizing Committee**, Joint Workshop on High-Performance and Distributed Computing for Medical Imaging (HP-MICCAI/MICCAI-DCI) 2011
- * **Organizing Committee**, High-Performance MICCAI 2010
Jointly with Chaudhary (UB), Gong (IBM), Blezek (Mayo), and Kulikowski (Rutgers)
- * **Organizing Committee**, High-Performance MICCAI 2008
Jointly with Chaudhary (UB) and Gong (IBM)

PROGRAM COMMITTEE AND CONFERENCE REVIEWER

- * European Conference on Computer Vision
- * IEEE Conference on Computer Vision and Pattern Recognition
- * IEEE Conference on Robotics and Automation
- * IEEE/RSJ International Conference on Intelligent Robots and Systems
- * IEEE International Conference on Computer Vision
- * International Conference on Multimedia and Expo
- * Medical Image Computing and Computer Aided Intervention.
Above, years are not given for readability; service has been steady for most of these since 2006.
- * IEEE International Conference on Semantic Computing 2015
- * IEEE International Conference on Face and Gesture 2015
- * Graphical Models in Computer Vision (GMCV at ECCV 2014) 2014
- * IEEE International Workshop on Web-scale Vision and Social Media (VSM at CVPR 2014) 2014
- * IEEE International Conference on Semantic Computing (ICSC) 2014
- * International Workshop on Large Scale Visual Commerce (LSVisCom at ICCV) 2013
- * Workshop on Large-Scale Video Search and Mining (LSVSM'13 at ICCV) 2013
- * International Workshop on Action Recognition with a Large Number of Classes (at ICCV) 2013
- * ISPRS Workshop on Image Sequence Analysis (ISA) 2013
- * IEEE International Conference on Semantic Computing (ICSC) 2013
- * IEEE Workshop on Mobile Vision (at ICCV) 2013
- * IEEE Face and Gesture (FG) 2013
- * Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR) 2013
- * Fully 3D Image Reconstruction in Radiology and Nuclear Medicine Meeting 2013
- * Asian Conference on Computer Vision 2012
- * IEEE International Conference on Semantic Computing (ICSC) 2012
- * International Symposium on Visual Computing (ISVC) 2012
- * International Conference on Image Processing (ICIP) 2012
- * IEEE International Conference on Semantic Computing (ICSC) 2011
- * Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR) 2011
- * Int'l Workshop on Stochastic Image Grammars (SIG-11, at ICCV) 2011
- * IEEE International Conference on Semantic Computing (ICSC) 2010
- * IEEE Workshop on Mobile Vision (at ICCV) 2011
- * IEEE Workshop on Mobile Vision (at CVPR) 2010
- * Workshop on Probabilistic, Models for Medical Image Analysis (at MICCAI) 2009
- * Int'l Workshop on Stochastic Image Grammars (SIG-09, at CVPR 2009) 2009
- * Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR) 2009
- * IEEE 8th Intl Symposium on Signal Processing and Information Technology (ISSPIT) 2008
- * Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR) 2007

JOURNAL REVIEWER

- * ACM Computing Reviews
- * Artificial Intelligence
- * Computer Methods and Programs in Biomedicine
- * Computer Vision and Image Understanding
- * IEEE Multimedia
- * IEEE Signal Processing Letters
- * IEEE Transactions on Biomedical Engineering
- * IEEE Transactions on Image Processing
- * IEEE Transactions on Information Technology in Biomedicine
- * IEEE Transactions on Knowledge and Data Engineering
- * IEEE Transactions on Medical Imaging
- * IEEE Transactions on Multimedia
- * IEEE Transactions on Pattern Analysis and Machine Intelligence
- * IEEE Transactions on Systems, Man, and Cybernetics
- * Image and Vision Computing
- * International Journal of Computer Assisted Radiology and Surgery
- * International Journal of Computer Vision
- * Journal of Mathematical Imaging and Vision
- * Machine Vision and Applications
- * Medical Physics
- * Neuro Image
- * Pattern Recognition
- * Pattern Recognition Letters
- * Robotics and Autonomous Systems

PANELIST

- * National Institutes of Health 2015–16
- * National Science Foundation Panel/Reviewer (CISE/IIS) 2009, 2011–13, 2015–21, 2023
- * Technology Foundation STW 2009
A Dutch funding agency for academic research in the field of applied sciences

UNIVERSITY OF MICHIGAN COMMITTEES

- * ECE Strategic Vision Committee 1920
- * CoE Robotics Institute Executive Committee 1617–1819
- * UM Institutional Autonomous Systems Committee 1617–1718
- * MTC Internal Funding Review Panel 1516
- * ECE Graduate Advisor for Computer Vision 1516–Current
- * CE Undergraduate Program Committee 1516–Current
- * CE Undergraduate Advisor 1516–1617
- * CE Undergraduate Honors and Awards 1516–1819
- * MTC Internal Funding Review Panel 1415
- * ECE Graduate Admissions Committee 1415–Current
- * CoE TRW Endowed Research Award Committee 1415
- * ECE Graduate Activities Committee 1415

SUNY AT BUFFALO COMMITTEES

- * Rising Scholar Committee (advisor to VPR) 2013–2014
- * Faculty Search Committee 2010–2014
- * Graduate Studies Committee 2008–2014
- * Artificial Intelligence Area Coordinator 2011–2013
- * Distinguished Speaker Series Committee 2010
- * Brochure and Website Committee 2010
- * Colloquium Committee 2009
- * Graduate Admissions Committee 2008–2009

INVITED TALKS

Curious Embeddings, Hazy Oracles, and the Path to Safe, Cooperative AI

Curated version of this talk is at YouTube: <https://youtu.be/Ck3ma6L4lug>

- SUNY at Buffalo CSE Seminar Series 10/2024
- Oregon State University CS Seminar Series 11/2024
- University of Pennsylvania ASSET Center Series 12/2024
- Johns Hopkins University CS Distinguished Speaker Series 12/2024

To Infer Or To Defer: Hazy Oracles in Human+AI Collaboration	
- Loyola University Maryland	9/2024
- IEEE CVPR Workshop on Computer Vision with Humans in the Loop	6/2024
- University of Rochester	3/2024
Everyday Depth: Towards Embodied AI with Depth from Standard Video Processing	
- (Keynote) IEEE CVPR Workshop on Visual Odometry and Computer Vision	6/2024
Video Understanding in the Clinic: Progress and Challenges	
- Michigan Institute for Data Science	9/2023
- (Keynote) AAAI Fall Symposium	11/2021
Me AI; You Human: Advances in Human-AI Cooperation	
- CORL Workshop on Bridging the Gap	11/2023
- IROS Workshop on Extended Reality	10/2022
- Purdue University	12/2021
- University of Michigan	10/2021
- Stevens Institute of Technology	2/2020
Rethinking Human Guidance in Video Object Segmentation	
- ICCV 2021 Workshop on Open-World Video Object Detection and Segmentation Challenge	10/2021
The Evolving Human-AI Relationship and its Impact on Education	
- American Society of Engineering Education	4/2021
Toward Kinematic Reconstruction of Roadway Scenes from Monocular Video using Hybrid Intelligence	
- University of Bonn Computer Science	11/2019
- University of Michigan Automotive Research Center Symposium	9/2018
Are we <i>using</i> our humans best? A View From Video Object Segmentation	
- British Computer Society Video Understanding Symposium	9/2019
Academic Perspectives for Video Analytics in Public Safety	
- NIST/DHS VAPS	11/2018
Segmentation is Selection	
- ECCV Workshop on Video Segmentation	9/2018
- NCSU ECE Distinguished Lecture Series	10/2018
Computer Vision & Video Analytics – Bridging Academia and Startup	
- ACM SIGIR	7/2018
Computer Vision & Video Analytics in Large-Scale Public Safety	
- Disney Large Venue Public Safety Meeting	8/2017
Toward Automated Understanding of Instructional Videos	
- INRIA Paris	4/2017
- University of Trento	4/2017
- ECCV Workshop on Visual Storytelling and Large Scale Movie Description Challenge	10/2016
Sparse Modeling in Deep and Cross-Modal Embeddings	
- Carnegie Mellon University LTI Colloquium	1/2017
- Information Theory and Applications	1/2017
Grouping Process Models in Actor-Action Segmentation	
- ActivityNet Large Scale Activity Recognition Challenge at IEEE CVPR	6/2016
- Xerox PARC	5/2016
- INRIA Paris	5/2016
Cross-Modal Embeddings with Video, Text and Speech	
- University of Rochester, Data Science Institute	5/2016
- Rochester Institute of Technology, Center for Imaging Science	5/2016
- Information Theory and Applications	1/2016
- Large Scale Movie Description Challenge Workshop at IEEE ICCV	12/2015
- Language and Vision Workshop at IEEE CVPR	6/2015
Video Analysis for Body Worn Cameras	
- (Keynote) Moving Cameras Meet Video Surveillance Workshop at IEEE CVPR 2016	6/2016
- Computers, Freedom, Privacy	10/2015
Supervoxels in Multi-X Understanding of Dynamic Scenes	

- (Keynote) IEEE/ISPRS 3rd Joint Workshop on Multi-Sensor Fusion for Dynamic Scene Understanding at IEEE CVPR	6/2015
Toward the Who and Where of Action Recognition	
- University of Maryland at College Park	9/2015
- INRIA Grenoble	7/2015
Perceiving Action in Space-Time: Computational and Human Perspectives	
- Michigan State University (CSE)	11/2014
- University of Michigan (EECS)	3/2014
- University of Minnesota (CSE)	3/2014
- Northeastern University (ECE)	2/2014
Why Label When You Can Compare? Active Constraint Pursuit in Metric Learning and Clustering	
- Virginia Tech (CS)	3/2014
Semantic Scale Selection from Video Segmentation Hierarchies	
- Ohio State University (ECE)	1/2014
What Representation is Right for Recognizing Activities in the Large Scale?	
- Keynote Address at ICCV 2013 Workshop on Action Recognition with a Large Number of Classes	12/2013
Can Language and Segmentation Play a Role in Large Scale Video Search?	
- ICCV 2013 Workshop on Large Scale Video Search and Mining	12/2013
Advances in Segmentation for Video Understanding	
- Rutgers University (Perceptual Science Forum Keynote Address)	5/2013
- University of Central Florida	4/2013
- University of Massachusetts Amherst	2/2013
- Stevens Institute of Technology	12/2012
- University of Pennsylvania	11/2012
- Microsoft Research Asia	8/2012
Joint Segmentation and Recognition of Medical Images with Layered Models	
- Second Vision and Multimedia Meeting in the Greater NY Area	6/2012
Graph-Shifts: Dynamic Hierarchical Energy Minimization for Segmentation and Classification.	
- GE Research	10/2010
- University of California, San Diego, Computer Science and Engineering	9/2010
- Army Research Labs CISD	10/2010
- Johns Hopkins University, Computer Science	5/2009
- Siemens Medical Solutions	12/2008
- Massachusetts Institute of Technology, CSAIL	11/2008
- SUNY at Buffalo, Computer Science and Engineering	3/2007
Bayesian Machine Intelligence Research at Buffalo	
- ITT Space Systems Division	2/2009
- MTM Interactive	6/2008
Multilevel Image Segmentation and Integrated Bayesian Model Classification.	
- Stony Brook University, Computer Science	10/2006
- IBM TJ Watson Research Center	10/2006
- Rutgers University, Center for BioImaging and Modeling	10/2006
- Siemens Corporate Research	10/2006
- Vanderbilt Institute for Imaging Science	7/2006
- UCLA Laboratory of Neuroimaging SIG-STAT Meeting	3/2006
- UCLA Statistics Seminar Series	4/2006
- Johns Hopkins University, ERC CISST	5/2006
Coherent Image Regions	
- UCLA Medical Imaging Informatics	5/2005

TEACHING

OUTREACH AND K12 TEACHING	
Exploring Careers in AI	4/2024
Nichols School 8th Grade class.	
AI Family Challenge	9/2018
Panel participant and mentor.	
Lights, Pinhole and Lenses	6/2016, 6/2017 & 6/2018

- U Michigan COE Xplore Program
Reached 60 Middle Schoolers through 6 90 minute sessions over two days.
- Lights, Pinhole and Lenses: Computer Imaging** 6/2015
U Michigan COE Xplore Program
Reached 60 Middle Schoolers through 6 90 minute sessions over two days.
- TUTORIALS AND SUMMER SCHOOLS
- Machine Learning II (Ensemble Methods) 6/2015
U Michigan Big Data Summer Institute Lecture
- Video Segmentation Tutorial at CVPR 2014 6/2014
with Irfan Essa and Matthias Grundmann
- IPAM Graduate Summer School on Computer Vision 7/2013
University of California, Los Angeles, Institute for Pure and Applied Mathematics
- International Summer School on Vision, Learning and Cognition 8/2012
Beijing University of Posts and Telecommunications
- UNIVERSITY OF MICHIGAN, ELECTRICAL ENGINEERING AND COMPUTER SCIENCE
- ROB 498/599 Special Topics: 3D Robot Perception**
- Fall 2023: Enrollment: 25
- EECS 542 Advanced Topics in Computer Vision**
- Winter 2017: Enrollment 75
- Winter 2016: Enrollment 18
- EECS 504 Foundations of Computer Vision**
- Fall 2023: Enrollment: 121
- Fall 2020: Enrollment: 130
- Fall 2018: Enrollment: 67
- Winter 2018: Enrollment 78
- Fall 2016: Enrollment 64
- Fall 2015 (as 598): Enrollment 41
- Fall 2014 (as 598): Enrollment 46
- EECS 598 Probabilistic Graphical Models**
- Winter 2015: Enrollment 36
- EECS 442 Computer Vision**
- Fall 2018: Enrollment 111
- Fall 2017: Enrollment 89
- SUNY AT BUFFALO, COMPUTER SCIENCE AND ENGINEERING
- CSE 455/555 Introduction to Pattern Recognition**
Cross-listed undergraduate graduate course.
- Spring 2013: Enrollment 70
- Spring 2012: Enrollment 73
- Spring 2011: Enrollment 60
- Spring 2010: Enrollment 35
- Spring 2009: Enrollment 29
- CSE 672 Bayesian Vision**
New upper-level graduate course.
- Fall 2012: Enrollment 22
- Fall 2010: Enrollment 12
- Spring 2009: Enrollment 8
- CSE 642 Techniques in AI: Vision for HCI**
- Fall 2009: Enrollment 8
- CSE Seminars**
- Spring 2014: Readings in Joint Visual, Lingual and Physical Models and Inference
- Fall 2011: Readings in Computer Vision and Machine Learning
- Fall 2010: Readings in Video Analysis
- Fall 2009: Readings in Image Semantics
- Fall 2008: Readings in Pattern Theory
- Fall 2007: Readings in Medical Image Segmentation

PRIOR EXPERIENCE AS A POSTDOC AND GRADUATE STUDENT

Object-Oriented Methods in Software Engineering - UCLA BIOMED 223C, Spring 2006.

Research Experience for Undergraduates Mentorship - Luz Molinelli, Summer 2005.

Project: Retinal disease diagnosis through multidimensional histogram analysis.

Research Experience for Undergraduates Mentorship - Ravi Mody, Summer 2004.

Project: Machine learning to track hand postures.

Guest Lecturer - Computer Vision, Cameras and Calibration, Dr. Gregory Hager, Fall 2002.

Teaching Assistant - Computer Vision, Dr. Gregory Hager, Fall 2001. Duties included holding weekly office hours, grading homeworks and projects, and preparing review notes.

Teaching Assistant - Data Structures, Dr. Subodh Kumar in Fall 2000 and Dr. Jonathan Cohen in Spring 2001. Duties included managing a team of 5 course assistants, holding weekly office hours, grading projects and exams, and holding review sessions.

STUDENT ADVISING

PH.D. STUDENTS

University of Michigan

D11. Yayuan Li EXPECTED 2027

D10. Filipos Bellos EXPECTED 2027

D9. Stephan Lemmer (ROB) 6/2023

THESIS: *Hazy Oracles in Deep Learning*

INITIAL PLACEMENT: AI Engineer at AEM Corporation

D8. Nathan Louis (ECE) 6/2024

THESIS: *Improving Articulated Pose Tracking and Contact Force Estimation for Qualitative Assessment of Human Actions*

INITIAL PLACEMENT: Research Scientist at Southwest Research Institute

D7. Madan Ravi Ganesh (ECE) 6/2022

THESIS: *Compression and Curriculum Strategies for Efficient Learning in Deep Neural Networks*

INITIAL PLACEMENT: Research Scientist at Bosch AI

D6. Shurjo Banerjee (ECE) 1/2022

THESIS: *Towards Improving Deep Visio-Linguistic Navigation*

INITIAL PLACEMENT: Research Scientist at Kodiak Robotics

D5. Byungsu Kyle Min (ECE) 11/2021

THESIS: *Video Understanding with Minimal Human Supervision*

INITIAL PLACEMENT: AI/ML Research Scientist at Intel

D4. Ryan Szeto (CSE) 8/2021

THESIS: *Enforcing Realism and Temporal Consistency for Large-Scale Video Inpainting*

INITIAL PLACEMENT: Senior Computer Vision Engineer at SafelyYou

D3. Luowei Zhou (ROB) 6/2020

THESIS: *Language-Driven Video Understanding*

INITIAL PLACEMENT: Research Scientist at Microsoft

CURRENT POSITION: Research Scientist at Google Brain

D2. Vikas Dhiman (ECE) 6/2019

THESIS: *Towards Better Navigation*

INITIAL PLACEMENT: PostDoc at University of California San Diego

CURRENT POSITION: Assistant Professor at the University of Maine ECE

D1. Chenliang Xu (CSE) 8/2016

THESIS: *Scale-Adaptive Video Understanding*

INITIAL PLACEMENT: Assistant Professor at University of Rochester CS

CURRENT POSITION: Associate Professor at the University of Rochester CS

SUNY at Buffalo

D'9. Duygu Sarikaya 6/2017

THESIS: *Exploring Fusion Models in Computer Vision for Medical Imaging*

INITIAL PLACEMENT: Postdoctoral Fellow at MediCIS

CURRENT POSITION: Lecturer at the University of Leeds

D'8. David Johnson 6/2017

THESIS: *From Hierarchies to Metrics: Learning Nonlinear Models of Semantics Association*

INITIAL PLACEMENT: Postdoctoral Fellow at the University of Michigan

- D'7. Wei Chen 8/2015
 THESIS: *Video Based Action Understanding*
 INITIAL PLACEMENT: Software Engineer at Amazon
- D'6. Ran Xu 8/2015
 THESIS: *Structured Models for Video Understanding*
 INITIAL PLACEMENT: Data Scientist at Banjo
- D'5. Kevin R. Keane 8/2015
 THESIS: *Implementing High Dimensional Gaussian Models For Financial Applications*
 INITIAL PLACEMENT: Systems Trading Group, Maple Securities USA, Inc.
- D'4. Caiming Xiong 8/2014
 THESIS: *Learning From and Actively Selecting Pairwise Constraints in Data Science*
 INITIAL PLACEMENT: Postdoc at UCLA
 CURRENT: VP of AI Research and Applied AI at Salesforce
- D'3. Jeffrey A. Delmerico 8/2013
 THESIS: *Attributed Object Maps: Descriptive Object Models as High-level Semantic Features for Mobile Robotics*
 PLACEMENT: Postdoc at ETH Zurich
 CURRENT: Senior Scientist at Microsoft
- D'2. Albert Y. C. Chen 5/2013
 THESIS: *Modeling and Optimizing Spatiotemporal Priors for Video Analysis Problems*
 INITIAL PLACEMENT: Computer Scientist at Tandent Vision Science (Pittsburgh, PA)
 CURRENT: Applied Scientist Manager at Amazon
- D'1. Ifeoma Nwogu, Ph.D. 8/2009
 Jointly advised with Prof. Venu Govindaraju
 THESIS: *An Ontology Driven Probabilistic Methodology for Image Understanding*
 PLACEMENT: NSF Computing Innovation Fellowship at U. of Rochester
 CURRENT: Assistant Professor at SUNY at Buffalo

M.S.E. STUDENTS

University of Michigan

- M5. Wensong Hu (ROB) 2024
- M4. Wuao Liu (ROB) 2023
- M3. Parker Koch (ECE) 2022
- M2. Victoria Florence (ROB) 2021
- M1. Eric Hofesmann (CSE) 2020

SUNY at Buffalo

- M'8. Duygu Sarikaya 6/2012
 THESIS: *Detection and Segmentation of Free Blood in FAST Exam Ultrasound Images*
 PLACEMENT: Ms. Sarikaya is a Ph.D. student at SUNY at Buffalo
- M'7. Ananth Sadanand 2/2012
 THESIS: *Action Bank: A High-Level Representation of Activity in Video*
 INITIAL PLACEMENT: Mr. Sadanand is a Computer Vision Engineer at Liveclips.com in New York, NY.
 CURRENT POSITION: Senior Applied Scientist at Amazon
- M'6. Sagar Waghmare 2/2012
 THESIS: *Comparative Study of Feature-Selective Sliding Window Object Detectors in Images*
 PLACEMENT: Mr. Waghmare is a Computer Vision Programming Specialist at a startup in New York City.
- M'5. Xin Li, M.S.E. 8/2011
 THESIS: *Key-Part Detection Using Boundary-Regional Codebook*
 PLACEMENT: Mr. Li is a Ph.D. student at Temple University.
- M'4. Yingjie Miao, M.S.E. 6/2011
 PROJECT: *Hamiltonian Streamline Guided Feature Extraction with Application to Face Detection*
 PLACEMENT: Dr. Miao is a Machine Learning Engineer at SAP.
- M'3. Timothy J. Burns, M.S.E. 6/2010
 PROJECT: *Document Image Segmentation and Labeling with Topic Models*
 PLACEMENT: Mr. Burns is a Software Engineer in Buffalo, NY.

- M'2. Dipankar Das, M.S.E. 12/2009
 THESIS: *Hierarchical Multiple Instance Learning for Object Detection*
 PLACEMENT: Mr. Das is a Research Engineer at Ikona Medical in Santa Monica, CA.
- M'1. Chris Hoeflich, M.S.E. 12/2008
 PROJECT: *Segmentation of 2D Gel Electrophoresis Spots Using a Markov Random Field.*
 PLACEMENT: Mr. Hoeflich is a Software Engineer in Buffalo, NY.

UNDERGRADUATE STUDENTS

In the capacity that I mentor them in research.

University of Michigan

- U6. Max Ernst 6/2024
 U5. Nolan Kuza 6/2023
 U4. Samuel Tenka 6/2018
 U3. Max Morrison 6/2018
 U2. Yichen Yao 6/2017
 U1. Sajjan Patel 6/2016

SUNY at Buffalo

- U'6. Spencer Whitt 6/2014
 U'5. Philip Rosebrough 6/2012
 U'4. Steven Hsieh 6/2012
 U'3. Colin Lea 6/2011
 U'2. Alexander Haynie 6/2011
 U'1. Andrew Schlackman 6/2009

PH.D. AND M.S.E. COMMITTEE MEMBERSHIP

Students for whom I am not the primary advisor.

University of Michigan

- Z15. Shane Storks, Ph.D. Candidate (CSE) GRANTED SPRING 2023
 ADVISOR: Prof. Joyce Chai (CSE)
Coherent Physical Commonsense Reasoning in Foundational Language Models
- Z14. Xiaotong Chen, Ph.D. Candidate (ROB) GRANTED SPRING 2022
 ADVISOR: Prof. Chad Jenkins (ROB)
Affordance-grounded Robot Perception and Manipulation in Adversarial, Translucent, and Cluttered Environments
- Z13. Xinchun Yan, Ph.D. Candidate (CSE) GRANTED SPRING 2019
 ADVISOR: Prof. Honglak Lee
Learning Deep Controllable and Structured Representations for Image Synthesis, Structure Prediction, and Beyond
- Z12. Alexander Kalinin, Ph.D. Candidate (Bioinformatics) GRANTED FALL 2018
 ADVISORS: Prof. Brian Athey and Prof. Ivo Dinov
3D Cell Nuclear Morphology Analysis
- Z11. Arash Ushani, Ph.D. Candidate (CSE) GRANTED SUMMER 2018
 ADVISOR: Prof. Ryan Eustice
Understanding a Dynamic World: Dynamic Motion Estimation for Autonomous Driving Using LIDAR
- Z10. Yu-Wei Chao, Ph.D. Candidate (CSE) GRANTED SPRING 2018
 ADVISORS: Prof. Jia Deng
Visual Recognition and Synthesis of Human-Object Interactions
- Z9. John Lipor, Ph.D. (EE:Systems) GRANTED FALL 2017
 ADVISORS: Prof. Laura Balzano
Sensing Structured Signals with Active and Ensemble Methods
- Z8. Jie Li, Ph.D. (EE:Systems) GRANTED SUMMER 2017
 ADVISORS: Prof. Matthew Johnson-Roberson and Prof. Ryan Eustice
Place Recognition and Localization for Multi-Modal Underwater Navigation with Vision and Acoustic Sensors

- Z7. Matthew Jacobs, Ph.D. (Math) GRANTED SPRING 2017
 ADVISORS: Prof. Selim Esedoğlu and Jinho Baik
Algorithms for Multiphase Partitioning
- Z6. Suyog Jain, Ph.D (UT Austin) GRANTED SPRING 2017
 ADVISOR: Prof. Kristen Grauman (UT Austin)
Human Machine Collaboration for Foreground Segmentation in Images and Videos
- Z5. Guillaume Seguin, Ph.D. (INRIA Paris) GRANTED SPRING 2016
 ADVISOR: Prof. Ivan Laptev (INRIA Paris)
Analyse des personnes dans les films stéréoscopiques
- Z4. Yelin Kim, Ph.D (CSE) GRANTED SPRING 2016
 ADVISOR: Prof. Emily Mowers Provost
Automatic Emotion Recognition: Quantifying Dynamics and Structure in Human Behavior
- Z3. Dan Oneata, Ph.D. (INRIA Grenoble) GRANTED SUMMER 2015
 ADVISOR: Prof. Cordelia Schmid (INRIA Grenoble, U. Grenoble)
Modèles robustes et efficaces pour la reconnaissance d'actions et leur localisation (Robust and efficient models for action recognition and localization)
- Z2. Yu Xiang, Ph.D. (EE:Systems) GRANTED WINTER 2015
 ADVISORS: Prof. Silvio Savarese and Prof. Alfred Hero
3D Object Representations for Recognition
- Z1. Paul Ozog, Ph.D. (EE:Systems) GRANTED WINTER 2015
 ADVISOR: Prof. Ryan Eustice
Advances in Simultaneous Localization and Mapping in Confined Underwater Environments Using Sonar and Optical Imaging
- SUNY at Buffalo
- Z'12. Jingteng Xue, Ph.D. GRANTED SPRING 2014
 ADVISOR: Prof. Chang Wen Chen
Perceptual Quality Driven Video Evaluation and Processing
- Z'11. Wenyuan Yin, Ph.D. GRANTED SUMMER 2013
 ADVISOR: Prof. Chang Wen Chen
Mobile Multimedia: From Acquisition to Adaptation with Semantics, Context and Social Information
- Z'10. Shujie Liu, Ph.D. GRANTED SPRING 2013
 ADVISOR: Prof. Chang Wen Chen
Compression, Rendering and Transmission for 3D and Scalable Video
- Z'9. Yi Tang, Ph.D. GRANTED SUMMER 2012
 ADVISOR: Prof. Sargur Srihari
Evaluating the Probability of Identification in Forensic Science
- Z'8. Chang Su, Ph.D. GRANTED SUMMER 2011
 ADVISOR: Prof. Sargur Srihari
Machine Learning in Fingerprint Probability Evaluation
- Z'7. Ricardo Rodrigues, Ph.D. GRANTED SUMMER 2011
 ADVISOR: Prof. Venu Govindaraju
Face Modeling and Biometric Anti-Spoofing using Probability Distribution Transfer Learning
- Z'6. Anh Ngoc Le, Ph.D. GRANTED SPRING 2011
 ADVISOR: Prof. Hung Ngo.
On the Data Flow Masquerading Problem
- Z'5. Xujun Peng, Ph.D. GRANTED FALL 2010
 ADVISOR: Prof. Venu Govindaraju.
Probabilistic Random Field Based Text Identification From Annotated Machine Printed Documents
- Z'4. Gabriel Terajanu, Ph.D. GRANTED SPRING 2010
 ADVISOR: Prof. Peter Scott
Towards a Decision-Centric Framework for Uncertainty Propagation and Data Assimilation
- Z'3. Bhaskar Purkayastha, M.S.E. GRANTED FALL 2009
 ADVISOR: Prof. Venu Govindaraju
Integrating Gesture Recognition and Speech Recognition in a Touch-Less Human Computer Interaction System.
- Z'2. Peter Noël, Ph.D. GRANTED SUMMER 2009
 ADVISOR: Prof. Jinhui Xu
Geometric Algorithms for Three Dimensional Reconstruction in Medical Imaging

Z'1. Hani Z. Girgis, Ph.D.
ADVISOR: Prof. Daniel Fischer
Machine Learning Based Meta Approaches to Protein Structure Prediction

GRANTED SUMMER 2008

VISITING SCHOLARS

V8. Haonan Chen (Tsinghua University, China; Undergrad) 6/2016–9/2016
V7. Jiamei Yan (Zhejiang University, China; Undergrad) 6/2016–9/2016
V6. Baiyun Cui (Zhejiang University, China; Undergrad) 6/2016–9/2016
V5. Tingtin Han (Harbin Institute of Technology, China) 9/2015-8/2016
V4. David Molik (Rensselaer Polytechnic Institute, Undergrad) SUMMER 2012
V3. Wei Wu (Assoc. Professor, School of Information Engineering, Wuhan University of Technology) 2/2011-2/2012
V2. Srijan Kumar (IIT Kharagpur, Undergrad) SUMMER 2011
V1. Digvijay Gagneja (IIT Kharagpur, Undergrad) SUMMER 2010

INDUSTRIAL
ACTIVITIES

Co-Founder of and Operator at Voxel51, Inc ANN ARBOR, MI 12/2016–Current
Led the company as CEO from 2016–2023 closing both Seed (~\$2.5M) and Series A (~\$12M) funding rounds led by eLab Ventures and Drive Capital respectively, and growing the company to 20+ employees

Now Chief Scientist leading innovation for the next generation AI developer tooling.

Voxel51 is now a Series B company with Bessemer Ventures leading our ~\$30M latest round.

Board of Advisors for Reflective.AI SAN FRANCISCO, CA 6,2019–12/2023

Board of Advisors for Onu Technology, Inc. SAN FRANCISCO, CA 1,2017–12/2022

Consultant for SoarTech, Inc. ANN ARBOR, MI 3/2016–2019

Consultant for Toyota Research Institute, Inc. ANN ARBOR, MI 5/2016–4/2017

Consultant for Cyvek, Inc. WALLINGFORD, CT 4/2015–4/2016

Consultant for NEC Research, Inc. CUPERTINO, CA 8/2014–6/2015

Consultant for CapsoVision, Inc. SANTA CLARA, CA 8/2014–Current

Consultant for CUBRC, Inc. BUFFALO, NY 12/2009–3/2014
Various projects related to computer vision and image analysis for defense contracts.

Consultant for Ikona Medical Corp. LOS ANGELES, CA 9/2006–6/2014
Development of real-time medical video mosaicking algorithms and software.

Licensed Technology 9/2009
Licensed GPU based reconstruction algorithm for CT data to IRIS (Ionizing Radiation Imaging Systems LLC).

Co-Founder of NaviGuru.com LIVE 4/2006–11/2009
NaviGuru.com was a Web 2.0 site that unified social networking with on-line mapping technology. It introduced a new form of web search called *visual spatial query*.

Consultant for Infinite Biomedical Technologies, LLC BALTIMORE, MD 9/2006–9/2008
Development of image calibration and dewarping algorithm used in contact endoscopy.

Provisional Patent with Licensable Technology 2003

TITLE: 4D Touchpad - VICs based interface to computer systems.

INSTITUTE: Johns Hopkins University, Baltimore MD (JHU Ref. DM-4181).

CO-INVENTORS: Gregory D. Hager and Darius Burschka.

DESCRIPTION: The 4D Touchpad builds a shared perceptual space between the computer user and a set of video cameras. Perceptual gestures are used to directly interact with interface components. The video cameras sense and interpret the gestures and effect automation in the computer system.

FULL EDUCATION

University of California, Los Angeles LOS ANGELES, CA
Post-Doc in Neuroscience and Statistics 2006-2007
ADVISORS: Dr. Alan Yuille and Dr. Arthur Toga

The Johns Hopkins University BALTIMORE, MD
Ph.D. in Computer Science 6/2006

ADVISOR: Dr. Gregory D. Hager
DISSERTATION TITLE: *"Techniques for Vision-Based Human-Computer Interaction"*

The Johns Hopkins University BALTIMORE, MD
M.S.E. in Computer Science 6/2002
PROJECT 1 ADVISOR: Dr. Gregory D. Hager, Computational Interaction and Robotics Lab (CIRL)
PROJECT 1 TITLE: *"Planar Surface Tracking Using Direct Stereo"*
PROJECT 2 ADVISOR: Dr. Jonathan Cohen, Graphics Lab
PROJECT 2 TITLE: *"Out-Of-Core Voxelization of Large Scalar Fields for Interactive Multiresolution Volume Rendering"*

Loyola College in Maryland BALTIMORE, MD
B.S. in Computer Science, *Cum Laude, Ranked First in Major* 5/2000
ADVISOR: Dr. Roger Eastman

Chaminade High School MINEOLA, NY

PAST POSITIONS

Director of the Stevens Institute for Artificial Intelligence HOBOKEN, NJ
Viola Ward Brinning and Elbert Calhoun Brinning Professor of Computer Science (with tenure)
Stevens Institute of Technology 1/2021-1/2022

Associate Professor (with tenure) BUFFALO, NY
Computer Science and Engineering
State University of New York at Buffalo 8/2013 - 8/2014

Assistant Professor BUFFALO, NY
Computer Science and Engineering
State University of New York at Buffalo 8/2007-8/2013

Post-Doctoral Fellow LOS ANGELES, CA
Neuroscience and Statistics
University of California, Los Angeles 9/2006-11/2007

MENTORS: Drs. Alan Yuille and Arthur Toga
PRIMARY FOCUS: Develop automatic, efficient and robust segmentation and recognition techniques for computational neuroimaging problems with coupled statistical learning methods. Implement and deploy software tools based on these algorithms into the research community.

Post-Doctoral Fellow LOS ANGELES, CA
Radiological Sciences and Statistics
University of California, Los Angeles 9/2005-8/2006

MENTORS: Drs. Alan Yuille and Ricky Taira
PRIMARY FOCUS: Develop automatic segmentation and recognition techniques for medical imaging problems (e.g., brain tumor) by integrating bottom-up detection with top-down models. Quantify statistics of the models' shape and appearance to improve accuracy of diagnosis and treatment.

Research Assistant BALTIMORE, MD
Computer Science
The Johns Hopkins University 8/2001-8/2005
ADVISOR: Dr. Gregory D. Hager
PROJECT: Developing vision-based techniques enabling dynamic, complex interaction in immersive mixed-reality environments: the VICs project.

Research Intern PRINCETON, NJ
Siemens Corporate Research SUMMER 2003
MENTOR: Dr. Yakup Genc
PROJECT: Markerless, real-time camera pose tracking using stereo video for Augmented Reality.

Software Engineer BALTIMORE, MD
The Johns Hopkins University FALL 2001
DESCRIPTION: Contracted by the Department of Computer Science at The Johns Hopkins University to design and development a SQL-based database and WWW interface for the faculty recruitment/search process. The system remains in use today with no downtime.

Acting Director Of Technology BALTIMORE, MD
Bionic Box Inc. 5/2000-9/2000
DESCRIPTION: Responsible for all internal IT and managed all (participated in some) software development projects.

Software Engineer BALTIMORE, MD

Alexander and Tom, Inc. 9/1999-5/2000
URL: <http://www.alextom.com>
RESPONSIBILITIES: Design and development of a broad range of interactive systems including small video-games, database systems, websites, and custom interactive cd-roms.

Research Intern ROCKVILLE, MD
Earth Satellite Corporation 6/1999-12/2000
URL: <http://www.earthsat.com>
PROJECT: Modify and deploy NASA software for radiometrical and geometrical distortion correction for the Landsat 7 satellite.

Research Assistant BALTIMORE, MD
Computer Science
Loyola College in Maryland SPRING 1999
ADVISOR: Dr. Keith Gallagher
PROJECT: Development of an ISO 9000-3 compatible software project management tool.

Hauber Science Research Fellow BALTIMORE, MD
Computer Science
Loyola College in Maryland SUMMER 1998
ADVISOR: Dr. Roger Eastman
PROJECT: Development of an image-processing algorithm for robust registration of retinal nerve images for use in glaucoma diagnosis.

Database Programmer MANHATTAN, NY
Information Builders, Inc. SUMMER 1996, 1997 AND WINTER 1997
URL: <http://www.informationbuilders.com>
RESPONSIBILITIES: Fulfill internal database programming needs for information systems using their proprietary database language and development platforms (FOCUS).