

Vladlen Koltun
Curriculum Vitae

vladlen.info
vladlen.koltun@intel.com
Born Nov 1980

EDUCATION

- 2002 Ph.D., Computer Science, Tel Aviv University, with distinction
Thesis: "Arrangements in Four Dimensions and Related Structures"
Supervisor: Micha Sharir
- 2000 B.Sc., Computer Science, Tel Aviv University, Magna Cum Laude

EXPERIENCE

- 3/2018 – present Intel Labs, Santa Clara, CA
Senior Principal Researcher
Director, Intelligent Systems Lab
Renamed the lab to reflect its broad scope. The lab conducts high-impact basic research on intelligent systems, covering machine learning, robotics, computer vision, and graphics.
- 1/2015 – 2/2018 Intel Labs, Santa Clara, CA
Principal Researcher
Director, Visual Computing Lab
Built a research lab from scratch. Recruited and hired more than 20 research scientists, postdocs, and research engineers. Mentored more than 20 PhD interns. Starting with basic research in computer vision and graphics, expanded the scope of the lab to cover machine learning, sensorimotor control, and robotics. Created the Intel Network on Intelligent Systems, a network of collaborations with more than a dozen leading research labs in academia. Helped define a number of other strategic initiatives at Intel Labs. Established ongoing relationships with business units throughout Intel, resulting in multiple documented tech transfers. Expanded the lab internationally, establishing a European location in Munich and growing it to a critical mass of researchers, postdocs, and interns. Initiated a number of large-scale open-source projects, including Open3D and CARLA. Published high-impact research on deep network architectures, simulation for intelligent systems, image processing and synthesis, and sensorimotor control.
- 6/2013 – 12/2014 Adobe Research, San Francisco, CA
Senior Research Scientist
Conducted research in visual computing. Primary focus on three-dimensional reconstruction. Additional results in visual perception and motor control.
- 7/2005 – 12/2013 Computer Science Department, Stanford, CA
Assistant Professor
Joined the faculty as a theoretician. Theoretical research recognized with the National Science Foundation CAREER Award (2006) and the Sloan Research Fellowship (2007).

Switched to research in visual computing in 2007. Built a new research group. Raised funding, supervised PhD students and postdoctoral researchers, mentored more than 20 undergraduate and master's researchers, taught 16 courses. Two PhD students are now professors at top-10 computer science departments, UC Berkeley (Sergey Levine) and UT Austin (Philipp Krähenbühl). Work on data-driven 3D modeling was licensed and formed the basis of a widely-used character creation application (Mixamo Fuse, subsequently Adobe Fuse). Work on dense random fields was recognized with the highest award at NIPS, the leading machine learning conference. Was chosen as the "most influential Stanford academic advisor" by all three undergraduate researchers under mentorship who received the Terman Award for outstanding academic achievement.

- 8–12/2003 Mathematical Sciences Research Institute, Berkeley, CA
Research Fellow
Conducted research in theoretical computer science.
- 9/2002 – 6/2005 University of California, Berkeley, CA
Postdoctoral Researcher
Conducted research in theoretical computer science.
Supervisor: Christos Papadimitriou
- 6/1997 – 5/1999 Shells Interactive, Israel
Senior Software Developer

HONORS AND AWARDS

- Sloan Research Fellowship, Alfred P. Sloan Foundation, 2007
- NSF CAREER Award, National Science Foundation, 2006
- Presidential Grant for Junior Faculty, Stanford University, 2006
- David Morgenthaler II Faculty Scholarship, Stanford University, 2005
- Rothschild Postdoctoral Fellowship, 2002. Awarded annually to 12 graduate students from all scientific disciplines in Israel.
- Deutsch Prize, Tel Aviv University, 2001. Awarded annually to one graduate student "whose research was singled out as outstanding".
- Wolf Foundation Fellowship, 2001. Awarded annually to selected graduate students from all scientific disciplines in Israel.

PUBLICATION AWARDS AND DISTINCTIONS

- "*Acoustic Non-Line-of-Sight Imaging*"
Computer Vision and Pattern Recognition (CVPR), 2019.
Selected for oral presentation at the conference. Selection rate 5.6%. (288 papers selected out of 5160 submissions.)

- *“Deep Drone Racing: Learning Agile Flight in Dynamic Environments”*
Conference on Robot Learning (CoRL), 2018.
Best Systems Paper Award. 3 papers selected for awards of any kind, out of 237 submissions to the conference.
- *“Semi-parametric Image Synthesis”*
Computer Vision and Pattern Recognition (CVPR), 2018.
Selected for full oral presentation at the conference. Selection rate 2.1%. (70 papers selected out of 3303 submissions.)
- *“Photographic Image Synthesis with Cascaded Refinement Networks”*
International Conference on Computer Vision (ICCV), 2017.
Selected for full oral presentation at the conference. Selection rate 2.1%. (45 papers selected out of 2143 submissions.)
- *“Playing for Benchmarks”*
International Conference on Computer Vision (ICCV), 2017.
Selected for a spotlight oral presentation at the conference. Selection rate 4.7%. (101 papers selected for spotlight oral or full oral out of 2143 submissions.)
- *“Learning to Act by Predicting the Future”*
International Conference on Learning Representations (ICLR), 2017.
Selected for full oral presentation at the conference. Selection rate 3.0%. (15 papers selected out of 507 submissions.)
- *“Fast Global Registration”*
European Conference on Computer Vision (ECCV), 2016.
Selected for full oral presentation at the conference. Selection rate 1.8%. (28 papers selected out of 1561 submissions.)
- *“Full Flow: Optical Flow Estimation by Global Optimization over Regular Grids”*
Computer Vision and Pattern Recognition (CVPR), 2016.
Selected for full oral presentation at the conference. Selection rate 3.9%. (83 papers selected out of 2145 submissions.)
- *“Feature Space Optimization for Semantic Video Segmentation”*
Computer Vision and Pattern Recognition (CVPR), 2016.
Selected for full oral presentation at the conference. Selection rate 3.9%. (83 papers selected out of 2145 submissions.)
- *“Robust Nonrigid Registration by Convex Optimization”*
International Conference on Computer Vision (ICCV), 2015.
Selected for full oral presentation at the conference. Selection rate 3.3%. (56 papers selected out of 1698 submissions.)
- *“Geodesic Object Proposals”*
European Conference on Computer Vision (ECCV), 2014.
Selected for full oral presentation at the conference. Selection rate 2.8%. (38 papers selected out of 1359 submissions.)
- *“Elastic Fragments for Dense Scene Reconstruction”*
International Conference on Computer Vision (ICCV), 2013.
Selected for full oral presentation at the conference. Selection rate 2.5%. (41 papers selected out of 1629 submissions.)

- *“Efficient Inference in Fully Connected CRFs with Gaussian Edge Potentials”*
Advances in Neural Information Processing Systems (NIPS), 2011.
Outstanding Student Paper Award. Highest award at the conference. (There was no separate Best Paper award at the time.) 3 papers selected out of 1400 submissions to the conference.
- *“Metropolis Procedural Modeling”*
ACM Transactions on Graphics, 30(2), 2011.
Featured on the cover of ACM Transactions on Graphics.
- *“Computer-Generated Residential Building Layouts”*
ACM Transactions on Graphics, 29(6), 2010.
Featured on the title page of ACM Transactions on Graphics.
- *“On Overlays and Minimization Diagrams”*
22nd Symposium on Computational Geometry (SoCG), 2006.
Selected for special issue of the journal Discrete and Computational Geometry devoted to the conference.
6 papers selected out of 138 submissions to the conference.
- *“On the Union of κ -Round Objects in Three and Four Dimensions”*
20th ACM Symposium on Computational Geometry (SoCG), 2004.
Selected for special issue of the journal Discrete and Computational Geometry devoted to the conference.
8 papers selected out of 147 submissions to the conference.
- *“Polyhedral Voronoi Diagrams of Polyhedra in Three Dimensions”*
18th ACM Symposium on Computational Geometry (SoCG), 2002.
Selected for a special issue of the journal Discrete and Computational Geometry devoted to the conference. 7 papers selected out of 104 submissions to the conference.
- *“Almost Tight Upper Bounds for Vertical Decompositions in Four Dimensions”*
42nd IEEE Symposium on Foundations of Computer Science (FOCS), 2001.
Machtey Award for best student-authored paper. Highest award at the conference. (There was no separate Best Paper award at the time.) 2 papers selected out of 214 submissions to the conference.
- *“Segment Intersection Searching Problems in General Settings”*
17th ACM Symposium on Computational Geometry (SoCG), 2001.
Selected for special issue of the journal Discrete and Computational Geometry devoted to the conference.
8 papers selected out of 106 submissions to the conference.

KEYNOTES, DISTINGUISHED LECTURES, SELECTED INVITED TALKS

Not including departmental colloquia, regional meetings, research seminars, etc. Not including talks prior to 2010.

- Invited speaker, *International Image Sensor Workshop*. Snowbird, UT, 2019
- Keynote speaker, *CVPR Workshop: Autonomous Driving*. Long Beach, CA, 2019.
- Invited speaker, *CVPR Workshop: Deep Learning for Semantic Visual Navigation*. Long Beach, CA, 2019.
- Invited speaker, *CVPR Workshop: 3D Scene Understanding for Vision, Graphics, and Robotics*. Long Beach, CA, 2019.
- Invited speaker, *CVPR Workshop: 3D Scene Generation*. Long Beach, CA, 2019.

- Keynote speaker, *IEEE International Conference on Computational Photography (ICCP)*. Tokyo, Japan, 2019.
- Keynote speaker, *IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR)*. Taichung, Taiwan, 2018.
- Keynote speaker, *ECCV Workshop: Visual Learning and Embodied Agents in Simulation Environments*. Munich, Germany, 2018.
- Keynote speaker, *ECCV Workshop: Autonomous Navigation in Unconstrained Environments*. Munich, Germany, 2018.
- Keynote speaker, *ECCV Workshop: Video Segmentation and Representation in the Deep Learning Era*. Munich, Germany, 2018.
- Invited speaker, *Rank Prize Symposium on Geometry and Uncertainty in Deep Learning for Computer Vision*. Grasmere, England, 2018.
- Keynote speaker, *Symposium on Geometry Processing*. Paris, France, 2018.
- Invited speaker, *RSS Workshop: New Benchmarks, Metrics, and Competitions for Robotic Learning*. Pittsburgh, PA, 2018.
- Invited speaker, *CVPR Workshop: Real-World Challenges and New Benchmarks for Deep Learning in Robotic Vision*. Salt Lake City, UT, 2018.
- Invited speaker, *CVPR Workshop: Deep Learning for Visual SLAM*. Salt Lake City, UT, 2018.
- Invited speaker, *CVPR Workshop: How to be a good citizen of the CVPR community*. Salt Lake City, UT, 2018.
- Invited speaker, *CVPR Workshop: Autonomous Driving*. Salt Lake City, UT, 2018.
- Invited speaker, *Data Learning and Inference: Autonomous Driving*. Playa Blanca, Canary Islands, 2018.
- Invited speaker, *Scenes from Video III*. Lake Garda, Italy, 2017.
- Invited speaker, *ICCV Workshop: Beyond Supervised Learning*. Venice, Italy, 2017.
- Invited speaker, *ICCV Workshop: Joint COCO and Places Challenge*. Venice, Italy, 2017.
- Invited speaker, *ICCV Workshop: Role of Simulation in Computer Vision*. Venice, Italy, 2017.
- Invited speaker, *ICCV Workshop: Multiview Relationships in 3D Data*. Venice, Italy, 2017.
- Invited speaker, *CVPR Workshop: Bridges to 3D*. Honolulu, HI, 2017.
- Keynote speaker, *International Computer Vision Summer School*. Sicily, Italy, 2017.
- Invited speaker, *ECCV Workshop: Geometry Meets Deep Learning*. Amsterdam, The Netherlands, 2016.
- Invited speaker, *ECCV Workshop: Virtual/Augmented Reality for Visual Artificial Intelligence*. Amsterdam, The Netherlands, 2016.
- Invited speaker, *Scenes from Video II*. Colchagua Valley, Chile, 2015.
- Invited speaker, *ICCV Workshop: Inverse Rendering*. Santiago, Chile, 2015.
- Distinguished speaker, *Center for Vision and Virtual Reality, UNC. Center opening*. Chapel Hill, NC, 2015.
- Invited speaker, *CVPR Workshop: 3D from a Single Image*. Boston, MA, 2015.
- Invited speaker, *ECCV Workshop: Reconstruction Meets Recognition Challenge*. Zürich, Switzerland, 2014.

- Keynote speaker, *International Conference on 3D Vision (3DV)*. Seattle, WA, 2013.
- Keynote speaker, *Computational Visual Media Conference (CVM)*. Beijing, China, 2012.

SELECTED PROFESSIONAL ACTIVITIES

Not including program committee memberships, area chair roles, etc.

- Editorial board, *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2016 – present
- Program Chair, *Symposium on Computer Animation (SCA)*, 2014
- Guest Editor, *SIAM Journal on Computing*, Special Section on ACM Symposium on Theory of Computing, 2009
- Guest Editor, *SIAM Journal on Computing*, Special Section on Foundations of Computer Science, 2007
- Editorial board, *Theory of Computing*, 2005 – 2014

ADVISING

Research Scientists

- Dr. Matthias Müller. Intel Labs. 2019 – present.
- Dr. Sohil Shah. Intel Labs. 2018 – present.
- Dr. German Ros. Intel Labs. 2018 – present.
- Dr. Stephan Richter. Intel Labs. 2018 – present.
- Dr. Qifeng Chen. Intel Labs. 2017 – 2018. (Next: Faculty at HKUST.)
- Dr. Katelyn Gao. Intel Labs. 2017 – present.
- Dr. René Ranftl. Intel Labs. 2017 – present.
- Dr. Alexey Dosovitskiy. Intel Labs. 2017 – 2019. (Next: Google Brain.)
- Dr. Philipp Krähenbühl. Intel Labs. 2016 – present. (Part time.)
- Dr. Jaesik Park. Intel Labs. 2015 – 2019. (Next: Faculty at POSTECH.)
- Dr. Jia Xu. Intel Labs. 2015 – 2017. (Next: Principal Researcher, Tencent AI Lab, China.)
- Dr. Vibhav Vineet. Intel Labs. 2015 – 2017. (Next: Returned to England for family reasons.)
- Dr. Qian-Yi Zhou. Intel Labs. 2015 – 2017. (Next: Co-founded a start-up with experienced entrepreneurs.)

Postdoctoral Researchers

- Dr. Benjamin Ummerhofer. Intel Labs. 2018 – present.
- Dr. Ozan Sener. Intel Labs. 2018 – present.
- Dr. Gernot Riegler. Intel Labs. 2018 – present.
- Dr. Gil Ben-Artzi. Intel Labs. 2017 – 2018.

- Dr. Zhuwen Li. Intel Labs. 2017 – 2018.
- Dr. Alexey Dosovitskiy. Intel Labs. 2016.
- Dr. Arno Knapitsch. Intel Labs. 2015 – 2017.
- Dr. René Ranftl. Intel Labs. 2015 – 2016.
- Dr. Vibhav Vineet. Stanford University. 2014 – 2015.
- Dr. Qian-Yi Zhou. Stanford University. 2012 – 2015.
- Dr. Martin Bokeloh. Stanford University. 2011 – 2013.
- Dr. Jack Wang. Stanford University. 2010 – 2013.
- Dr. Evangelos Kalogerakis. Stanford University. 2010 – 2012.
- Dr. Paul Merrell. Stanford University. 2009 – 2011.

Doctoral Students

- Qifeng Chen, Computer Science, Stanford University. 2013 – 2017. (After PhD: Research Scientist at Intel Labs, then faculty at HKUST.)
- Sungjoon Choi, Electrical Engineering, Stanford University. 2012 – 2015. (After PhD: Google.)
- Philipp Krähenbühl, Computer Science, Stanford University. 2009 – 2014. (After PhD: Postdoc at UC Berkeley, then faculty at UT Austin.)
- Sergey Levine, Computer Science, Stanford University. 2009 – 2014. (After PhD: Postdoc at UC Berkeley, then faculty at UC Berkeley.)
- Siddhartha Chaudhuri, Computer Science, Stanford University. 2005 – 2011. (After PhD: Postdoc at Princeton, then lecturer at Cornell.)

PhD Committees (Partial list. Not including PhD committees prior to 2018.)

- Jemin Hwangbo, 2018, ETH Zürich. Supervisor: Marco Hutter.
- Brandon Amos, 2018, CMU. Supervisor: Zico Kolter.
- Chen Chen, 2018, UIUC. Supervisor: Minh Do.
- Pierre Baqué, 2018, EPFL. Supervisors: Pascal Fua and François Fleuret. (After PhD: CEO of Neural Concept)
- Abhijit Kundu, 2018, Georgia Tech. Supervisors: Jim Rehg and Frank Dellaert. (After PhD: Google Research)
- Fisher Yu, 2018, Princeton. Supervisor: Tom Funkhouser. (After PhD: Postdoc at UC Berkeley)

Doctoral Interns

- Johan Björck. Intel Labs, visiting from Cornell. Summer 2019.
- Wei Dong. Intel Labs, visiting from CMU. Summer 2019.
- Hexiang (Frank) Hu. Intel Labs, visiting from USC. Summer 2019.
- John Lambert. Intel Labs, visiting from Georgia Tech. Summer 2019.
- Zhuang Liu. Intel Labs, visiting from UC Berkeley. Summer 2019.

- Aleksei Petrenko. Intel Labs, visiting from USC. Summer 2019.
- Anirudh Vemula. Intel Labs, visiting from CMU. Summer 2019.
- Xingyi Zhou. Intel Labs, visiting from UT Austin. Summer 2019.
- Christopher Choy. Intel Labs, visiting from Stanford. Winter/spring 2019.
- Hassan Abu AlHaija. Intel Labs, visiting from University of Heidelberg. Winter/spring 2019.
- Hengshuang Zhao. Intel Labs, visiting from CUHK. Winter/spring 2019.
- Zhipeng Cai. Intel Labs, visiting from University of Adelaide. Winter/spring 2019.
- Richard Shin. Intel Labs, visiting from UC Berkeley. Fall/winter 2018/2019.
- Brandon Amos. Intel Labs, visiting from CMU. Summer/fall 2018.
- Shaojie Bai. Intel Labs, visiting from CMU. Summer/fall 2018.
- Adel Bibi. Intel Labs, visiting from KAUST. Summer/fall 2018.
- Yang Gao. Intel Labs, visiting from UC Berkeley. Summer/fall 2018.
- Katrin Lasinger. Intel Labs, visiting from ETH Zürich. Summer/fall 2018.
- David Lindell. Intel Labs, visiting from Stanford. Summer/fall 2018.
- Dmytro Mishkin. Intel Labs, visiting from Czech Technical University. Summer/fall 2018.
- Henri Rebecq. Intel Labs, visiting from University of Zürich. Summer/fall 2018.
- Abhay Yadav. Intel Labs, visiting from University of Maryland. Summer/fall 2018.
- Xuaner Zhang. Intel Labs, visiting from UC Berkeley. Summer/fall 2018.
- Matthias Vestner. Intel Labs, visiting from TU München. October 2017 – April 2018
- Nikolay Savinov. Intel Labs, visiting from ETH Zürich. June-December 2017
- Maxim Tatarchenko. Intel Labs, visiting from Freiburg. May-November 2017
- Sergey Zagoruyko. Intel Labs, visiting from ParisTech. September 2017 – February 2018. (After PhD: Postdoc at Facebook AI Research.)
- Xiaojuan Qi. Intel Labs, visiting from CUHK. May-November 2017. (After PhD: Postdoc at the University of Oxford.)
- François Germain. Intel Labs, visiting from Stanford. June 2017 – March 2018
- Matthias Müller. Intel Labs, visiting from KAUST. June 2017 – January 2018. (After PhD: Intel Labs.)
- Felipe Codevilla. Intel Labs, visiting from Barcelona. June-December 2017
- Chen Chen. Intel Labs, visiting from UIUC. June-December 2017. Second internship May-November 2018. (After PhD: Apple.)
- Stephan Richter. Intel Labs, visiting from TU Darmstadt. May-November 2015. Second internship November 2016 – May 2017. (After PhD: Intel Labs.)
- Zeeshan Hayder. Intel Labs, visiting from Australian National University. November 2016 – May 2017.
- Qifeng Chen. Intel Labs, visiting from Stanford. September 2016 – June 2017. (After PhD: Intel Labs.)

- Katelyn Gao. Intel Labs, visiting from Stanford. June-December 2016. (After PhD: Intel Labs.)
- Sohil Atul Shah. Intel Labs, visiting from University of Maryland. May-December 2016. (After PhD: Intel Labs.)
- Marc Khoury. Intel Labs, visiting from UC Berkeley. May-December 2016.
- Mohammad Haris Baig. Intel Labs, visiting from Dartmouth. June-December 2015. (After PhD: Apple.)
- Linguang Zhang. Intel Labs, visiting from Princeton. June-September 2015.
- Jakob Engel. Intel Labs, visiting from TU München. May-November 2015. (After PhD: Oculus Research.)
- Abhijit Kundu. Intel Labs, visiting from Georgia Tech. May-November 2015. (After PhD: Google Research)
- Fisher Yu. Intel Labs, visiting from Princeton. February-November 2015. (After PhD: Postdoc at UC Berkeley.)
- Igor Mordatch. Stanford, visiting from University of Washington. September 2012 – March 2013. (After PhD: Postdoc at UC Berkeley, then faculty at CMU.)

PUBLICATIONS

Conference papers that also appeared in journals are listed only as journal papers. For publications in computational geometry (prior to 2009), the author list is in alphabetical order.

Refereed conference proceedings (excluding papers that also appeared in journals):

1. David Lindell, Gordon Wetzstein, and Vladlen Koltun, "Acoustic Non-Line-of-Sight Imaging," *Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, June 2019.
2. Xuaner Zhang, Qifeng Chen, Ren Ng, and Vladlen Koltun, "Zoom to Learn, Learn to Zoom," *Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, June 2019.
3. Maxim Tatarchenko, Stephan Richter, René Ranftl, Zhuwen Li, Vladlen Koltun, and Thomas Brox, "What Do Single-view 3D Reconstruction Networks Learn?," *Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, June 2019.
4. Henri Rebecq, René Ranftl, Vladlen Koltun, and Davide Scaramuzza, "Events-to-Video: Bringing Modern Computer Vision to Event Cameras," *Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, June 2019.
5. Gernot Riegler, Yiyi Liao, Simon Donné, Vladlen Koltun, and Andreas Geiger, "Connecting the Dots: Learning Representations for Active Monocular Depth Estimation," *Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, June 2019.
6. Elia Kaufmann, Mathias Gehrig, Philipp Foehn, René Ranftl, Alexey Dosovitskiy, Vladlen Koltun, and Davide Scaramuzza, "Beauty and the Beast: Optimal Methods Meet Learning for Drone Racing," *International Conference on Robotics and Automation (ICRA)*, Montréal, Canada, May 2019.
7. Shaojie Bai, J. Zico Kolter, and Vladlen Koltun, "Trellis Networks for Sequence Modeling," *International*

Conference on Learning Representations (ICLR), New Orleans, Louisiana, May 2019.

8. Adel Bibi, Bernard Ghanem, Vladlen Koltun, and René Ranftl, “Deep Layers as Stochastic Solvers,” *International Conference on Learning Representations (ICLR)*, New Orleans, Louisiana, May 2019.
9. Ozan Sener and Vladlen Koltun, “Multi-Task Learning as Multi-Objective Optimization,” *Advances in Neural Information Processing Systems (NIPS)*, Montréal, Canada, December 2018.
10. Zhuwen Li, Qifeng Chen, and Vladlen Koltun, “Combinatorial Optimization with Graph Convolutional Networks and Guided Tree Search,” *Advances in Neural Information Processing Systems (NIPS)*, Montréal, Canada, December 2018.
11. Elia Kaufmann, Antonio Loquercio, René Ranftl, Alexey Dosovitskiy, Vladlen Koltun, and Davide Scaramuzza, “Deep Drone Racing: Learning Agile Flight in Dynamic Environments,” *Conference on Robot Learning (CoRL)*, Zürich, Switzerland, October 2018.
12. Matthias Müller, Alexey Dosovitskiy, Bernard Ghanem, and Vladlen Koltun, “Driving Policy Transfer via Modularity and Abstraction,” *Conference on Robot Learning (CoRL)*, Zürich, Switzerland, October 2018.
13. Artemij Amiranashvili, Alexey Dosovitskiy, Vladlen Koltun, and Thomas Brox, “The Role of Motion Perception in Reinforcement Learning with Dynamic Objects,” *Conference on Robot Learning (CoRL)*, Zürich, Switzerland, October 2018.
14. René Ranftl and Vladlen Koltun, “Deep Fundamental Matrix Estimation,” *European Conference on Computer Vision (ECCV)*, Munich, Germany, September 2018.
15. Felipe Codevilla, Antonio López, Vladlen Koltun, and Alexey Dosovitskiy, “On Offline Evaluation of Vision-based Driving Models,” *European Conference on Computer Vision (ECCV)*, Munich, Germany, September 2018.
16. Xiaojuan Qi, Qifeng Chen, Jiaya Jia, and Vladlen Koltun, “Semi-parametric Image Synthesis,” *Computer Vision and Pattern Recognition (CVPR)*, Salt Lake City, UT, June 2018.
17. Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun, “Learning to See in the Dark,” *Computer Vision and Pattern Recognition (CVPR)*, Salt Lake City, UT, June 2018.
18. Maxim Tatarchenko, Jaesik Park, Vladlen Koltun, and Qian-Yi Zhou, “Tangent Convolutions for Dense Prediction in 3D,” *Computer Vision and Pattern Recognition (CVPR)*, Salt Lake City, UT, June 2018.
19. Zhuwen Li, Qifeng Chen, and Vladlen Koltun, “Interactive Image Segmentation with Latent Diversity,” *Computer Vision and Pattern Recognition (CVPR)*, Salt Lake City, UT, June 2018.
20. Nikolay Savinov, Alexey Dosovitskiy, and Vladlen Koltun, “Semi-parametric Topological Memory for Navigation,” *International Conference on Learning Representations (ICLR)*, Vancouver, Canada, May 2018.
21. Artemij Amiranashvili, Alexey Dosovitskiy, Vladlen Koltun, and Thomas Brox, “TD or not TD: Analyzing the Role of Temporal Differencing in Deep Reinforcement Learning,” *International Conference on Learning Representations (ICLR)*, Vancouver, Canada, May 2018.
22. Felipe Codevilla, Matthias Müller, Antonio López, Vladlen Koltun, and Alexey Dosovitskiy, “End-to-end Driving via Conditional Imitation Learning,” *International Conference on Robotics and Automation (ICRA)*,

Brisbane, Australia, May 2018.

23. Mohammad Haris Baig, Vladlen Koltun, and Lorenzo Torresani, "Learning to Inpaint for Image Compression," *Advances in Neural Information Processing Systems (NIPS)*, Long Beach, CA, December 2017.
24. Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun, "CARLA: An Open Urban Driving Simulator," *Conference on Robot Learning (CoRL)*, Mountain View, CA, November 2017.
25. Qifeng Chen and Vladlen Koltun, "Photographic Image Synthesis with Cascaded Refinement Networks," *International Conference on Computer Vision (ICCV)*, Venice, Italy, October 2017.
26. Stephan R. Richter, Zeeshan Hayder, and Vladlen Koltun, "Playing for Benchmarks," *International Conference on Computer Vision (ICCV)*, Venice, Italy, October 2017.
27. Qifeng Chen, Jia Xu, and Vladlen Koltun, "Fast Image Processing with Fully-Convolutional Networks," *International Conference on Computer Vision (ICCV)*, Venice, Italy, October 2017.
28. Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun, "Colored Point Cloud Registration Revisited," *International Conference on Computer Vision (ICCV)*, Venice, Italy, October 2017.
29. Marc Khoury, Qian-Yi Zhou, and Vladlen Koltun, "Learning Compact Geometric Features," *International Conference on Computer Vision (ICCV)*, Venice, Italy, October 2017.
30. Fisher Yu, Vladlen Koltun, and Thomas Funkhouser, "Dilated Residual Networks," *Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, July 2017.
31. Jia Xu, René Ranftl, and Vladlen Koltun, "Accurate Optical Flow via Direct Cost Volume Processing," *Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, July 2017.
32. Alexey Dosovitskiy and Vladlen Koltun, "Learning to Act by Predicting the Future," *International Conference on Learning Representations (ICLR)*, Toulon, France, April 2017.
33. Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun, "Fast Global Registration," *European Conference on Computer Vision (ECCV)*, Amsterdam, The Netherlands, October 2016.
34. Stephan Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun, "Playing for Data: Ground Truth from Computer Games," *European Conference on Computer Vision (ECCV)*, Amsterdam, The Netherlands, October 2016.
35. Qifeng Chen and Vladlen Koltun, "Full Flow: Optical Flow Estimation by Global Optimization over Regular Grids," *Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, June 2016.
36. Abhijit Kundu, Vibhav Vineet, and Vladlen Koltun, "Feature Space Optimization for Semantic Video Segmentation," *Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, June 2016.
37. René Ranftl, Vibhav Vineet, Qifeng Chen, and Vladlen Koltun, "Dense Monocular Depth Estimation in Complex Dynamic Scenes," *Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, June 2016.
38. Fisher Yu and Vladlen Koltun, "Multi-Scale Context Aggregation by Dilated Convolutions," *International Conference on Learning Representations (ICLR)*, San Juan, PR, May 2016.

39. Qifeng Chen and Vladlen Koltun, "Robust Nonrigid Registration by Convex Optimization," *International Conference on Computer Vision (ICCV)*, Santiago, Chile, December 2015.
40. Philipp Krähenbühl and Vladlen Koltun, "Learning to Propose Objects," *Computer Vision and Pattern Recognition (CVPR)*, Boston, MA, June 2015.
41. Sungjoon Choi, Qian-Yi Zhou, and Vladlen Koltun, "Robust Reconstruction of Indoor Scenes," *Computer Vision and Pattern Recognition (CVPR)*, Boston, MA, June 2015.
42. Qian-Yi Zhou and Vladlen Koltun, "Depth Camera Tracking with Contour Cues," *Computer Vision and Pattern Recognition (CVPR)*, Boston, MA, June 2015.
43. Philipp Krähenbühl and Vladlen Koltun, "Geodesic Object Proposals," *European Conference on Computer Vision (ECCV)*, Zürich, Switzerland, September 2014.
44. Sergey Levine and Vladlen Koltun, "Learning Complex Neural Network Policies with Trajectory Optimization," *International Conference on Machine Learning (ICML)*, Beijing, China, June 2014.
45. Qifeng Chen and Vladlen Koltun, "Fast MRF Optimization with Application to Depth Reconstruction," *Computer Vision and Pattern Recognition (CVPR)*, Columbus, OH, June 2014.
46. Qian-Yi Zhou and Vladlen Koltun, "Simultaneous Localization and Calibration: Self-Calibration of Consumer Depth Cameras," *Computer Vision and Pattern Recognition (CVPR)*, Columbus, OH, June 2014.
47. Sergey Levine and Vladlen Koltun, "Variational Guided Policy Search," *Advances in Neural Information Processing Systems (NIPS)*, Lake Tahoe, NV, December 2013.
48. Qifeng Chen and Vladlen Koltun, "A Simple Model for Intrinsic Image Decomposition with Depth Cues," *International Conference on Computer Vision (ICCV)*, Sydney, Australia, December 2013.
49. Qian-Yi Zhou, Stephen Miller, and Vladlen Koltun, "Elastic Fragments for Dense Scene Reconstruction," *International Conference on Computer Vision (ICCV)*, Sydney, Australia, December 2013.
50. Philipp Krähenbühl and Vladlen Koltun, "Parameter Learning and Convergent Inference for Dense Random Fields," *International Conference on Machine Learning (ICML)*, Atlanta, GA, June 2013.
51. Sergey Levine and Vladlen Koltun, "Guided Policy Search," *International Conference on Machine Learning (ICML)*, Atlanta, GA, June 2013.
52. Philipp Krähenbühl and Vladlen Koltun, "Efficient Nonlocal Regularization for Optical Flow," *European Conference on Computer Vision (ECCV)*, Florence, Italy, October 2012.
53. Sergey Levine and Vladlen Koltun, "Continuous Inverse Optimal Control with Locally Optimal Examples," *International Conference on Machine Learning (ICML)*, Edinburgh, Scotland, June 2012.
54. Young Min Kim, Jennifer Dolson, Mike Sokolsky, Vladlen Koltun, and Sebastian Thrun, "Interactive Acquisition of Residential Floor Plans," *International Conference on Robotics and Automation (ICRA)*, St. Paul, MN, May 2012.
55. Philipp Krähenbühl and Vladlen Koltun, "Efficient Inference in Fully Connected CRFs with Gaussian Edge Potentials," *Advances in Neural Information Processing Systems (NIPS)*, Granada, Spain, December 2011.

56. Sergey Levine, Zoran Popovic, and Vladlen Koltun, "Nonlinear Inverse Reinforcement Learning with Gaussian Processes," *Advances in Neural Information Processing Systems (NIPS)*, Granada, Spain, December 2011.
57. Sergey Levine, Zoran Popovic and Vladlen Koltun, "Feature Construction for Inverse Reinforcement Learning," *Advances in Neural Information Processing Systems (NIPS)*, Vancouver, Canada, December 2010.
58. Sarel Har-Peled and Vladlen Koltun, "Separability with Outliers," *16th International Symposium on Algorithms and Computation*, Sanya, China, December 2005.
59. Jason Hartline and Vladlen Koltun, "Near-Optimal Pricing in Near-Linear Time," *9th Workshop on Algorithms and Data Structures*, Waterloo, Canada, August 2005.
60. Vladlen Koltun, "Pianos are not Flat: Rigid Motion Planning in Three Dimensions," *16th ACM-SIAM Symposium on Discrete Algorithms*, Vancouver, BC, January 2005.
61. Quanfu Fan, Alon Efrat, Vladlen Koltun, Shankar Krishnan, and Suresh Venkatasubramanian, "Hardware-Assisted Natural Neighbor Interpolation," *7th Workshop on Algorithm Engineering and Experiments*, Vancouver, BC, January 2005.
62. Sarel Har-Peled, Vladlen Koltun, Dezhen Song, and Ken Goldberg, "Efficient Algorithms for Shared Camera Control," *19th ACM Symposium on Computational Geometry*, San Diego, CA, June 2003.
63. Vladlen Koltun and Micha Sharir, "On the Overlay of Envelopes in Four Dimensions," *13th ACM-SIAM Symposium on Discrete Algorithms*, San Francisco, CA, January 2002.
64. Vladlen Koltun, Yiorgos Chrysanthou, and Daniel Cohen-Or, "Hardware-Accelerated From-region Visibility Using a Dual Ray Space," *12th Eurographics Workshop on Rendering*, London, England, June 2001.
65. Vladlen Koltun, Yiorgos Chrysanthou, and Daniel Cohen-Or, "Virtual Occluders: an Efficient Intermediate PVS Representation," *11th Eurographics Workshop on Rendering*, Brno, Czech Republic, June 2000.

Refereed journals:

66. Delio Vicini, Vladlen Koltun, and Wenzel Jakob, "A Learned Shape-Adaptive Subsurface Scattering Model," *ACM Transactions on Graphics*, 38(4), 2019. Proceedings of ACM SIGGRAPH.
67. Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter, "Learning Agile and Dynamic Motor Skills for Legged Robots," *Science Robotics*, 4(26), 2019.
68. Jan Carius, René Ranftl, Vladlen Koltun, and Marco Hutter, "Trajectory Optimization with Implicit Hard Contacts," *IEEE Robotics and Automation Letters*, 3(4), 2018.
69. Jakob Engel, Vladlen Koltun, and Daniel Cremers, "Direct Sparse Odometry," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(3), 2018.
70. Sohil Atul Shah and Vladlen Koltun, "Robust Continuous Clustering," *Proceedings of the National Academy of Sciences (PNAS)*, 114(37), 2017.

71. Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun, "Tanks and Temples: Benchmarking Large-Scale Scene Reconstruction," *ACM Transactions on Graphics*, 36(4), 2017. Proceedings of ACM SIGGRAPH.
72. Qi-Xing Huang, Hai Wang, and Vladlen Koltun, "Single-View Reconstruction via Joint Analysis of Image and Shape Collections," *ACM Transactions on Graphics*, 34(4), 2015. Proceedings of ACM SIGGRAPH.
73. Qian-Yi Zhou and Vladlen Koltun, "Color Map Optimization for 3D Reconstruction with Consumer Depth Cameras," *ACM Transactions on Graphics*, 33(4), 2014. Proceedings of ACM SIGGRAPH.
74. Igor Mordatch, Jack Wang, Emanuel Todorov, and Vladlen Koltun, "Animating Human Lower Limbs Using Contact-Invariant Optimization," *ACM Transactions on Graphics*, 32(6), 2013. Proceedings of ACM SIGGRAPH Asia.
75. Qian-Yi Zhou and Vladlen Koltun, "Dense Scene Reconstruction with Points of Interest," *ACM Transactions on Graphics*, 32(4), 2013. Proceedings of ACM SIGGRAPH.
76. Jack Wang, Samuel Hamner, Scott Delp, and Vladlen Koltun, "Optimizing Locomotion Controllers Using Biologically-Based Actuators and Objectives," *ACM Transactions on Graphics*, 31(4), 2012. Proceedings of ACM SIGGRAPH.
77. Evangelos Kalogerakis, Siddhartha Chaudhuri, Daphne Koller, and Vladlen Koltun, "A Probabilistic Model for Component-Based Shape Synthesis," *ACM Transactions on Graphics*, 31(4), 2012. Proceedings of ACM SIGGRAPH.
78. Martin Bokeloh, Michael Wand, Hans-Peter Seidel, and Vladlen Koltun, "An Algebraic Model for Parameterized Shape Editing," *ACM Transactions on Graphics*, 31(4), 2012. Proceedings of ACM SIGGRAPH.
79. Sergey Levine, Jack Wang, Alexis Haraux, Zoran Popovic, and Vladlen Koltun, "Continuous Character Control with Low Dimensional Embeddings," *ACM Transactions on Graphics*, 31(4), 2012. Proceedings of ACM SIGGRAPH.
80. Martin Bokeloh, Michael Wand, Vladlen Koltun, and Hans-Peter Seidel, "Pattern-Aware Shape Deformation Using Sliding Dockers," *ACM Transactions on Graphics*, 30(6), 2011. Proceedings of ACM SIGGRAPH Asia.
81. Qi-Xing Huang, Vladlen Koltun, and Leonidas Guibas, "Joint Shape Segmentation with Linear Programming," *ACM Transactions on Graphics*, 30(6), 2011. Proceedings of ACM SIGGRAPH Asia.
82. Paul Merrell, Eric Schkufza, Zeyang Li, Maneesh Agrawala, and Vladlen Koltun, "Interactive Furniture Layout Using Interior Design Guidelines," *ACM Transactions on Graphics*, 30(4), 2011. Proceedings of ACM SIGGRAPH.
83. Siddhartha Chaudhuri, Evangelos Kalogerakis, Leonidas Guibas, and Vladlen Koltun, "Probabilistic Reasoning for Assembly-Based 3D Modeling," *ACM Transactions on Graphics*, 30(4), 2011. Proceedings of ACM SIGGRAPH.
84. Sergey Levine, Yongjoon Lee, Vladlen Koltun, and Zoran Popovic, "Space-Time Planning with Parameterized Locomotion Controllers," *ACM Transactions on Graphics*, 30(3), 2011. Presented at ACM SIGGRAPH.

85. Jerry O. Talton, Yu Lou, Steve Lesser, Jared Duke, Radomír Měch, and Vladlen Koltun, "Metropolis Procedural Modeling," *ACM Transactions on Graphics*, 30(2), 2011. Presented at ACM SIGGRAPH.
86. Siddhartha Chaudhuri and Vladlen Koltun, "Data-Driven Suggestions for Creativity Support in 3D Modeling," *ACM Transactions on Graphics*, 29(6), 2010. Proceedings of ACM SIGGRAPH Asia.
87. Paul Merrell, Eric Schkufza, and Vladlen Koltun, "Computer-Generated Residential Building Layouts," *ACM Transactions on Graphics*, 29(6), 2010. Proceedings of ACM SIGGRAPH Asia.
88. Sergey Levine, Philipp Krähenbühl, Sebastian Thrun, and Vladlen Koltun, "Gesture Controllers," *ACM Transactions on Graphics*, 29(4), 2010. Proceedings of ACM SIGGRAPH.
89. Sergey Levine, Christian Theobalt, and Vladlen Koltun, "Real-Time Prosody-Driven Synthesis of Body Language," *ACM Transactions on Graphics*, 28(5), 2009. Proceedings of ACM SIGGRAPH Asia.
90. Jerry O. Talton, Daniel Gibson, Lingfeng Yang, Pat Hanrahan, and Vladlen Koltun, "Exploratory Modeling with Collaborative Design Spaces," *ACM Transactions on Graphics*, 28(5), 2009. Proceedings of ACM SIGGRAPH Asia.
91. Siddhartha Chaudhuri and Vladlen Koltun, "Smoothed Analysis of Probabilistic Roadmaps," *Computational Geometry*, 42(8): 731–747, 2009.
92. Vladlen Koltun and Micha Sharir, "On Overlays and Minimization Diagrams," *Discrete and Computational Geometry*, 41(3): 385–397, 2009. Special issue on the 22nd Symposium on Computational Geometry.
93. Vladlen Koltun and Christos H. Papadimitriou, "Approximately Dominating Representatives." *Theoretical Computer Science*, 371, 148–154, 2007. Special issue on the 10th International Conference on Database Theory.
94. Boris Aronov, Alon Efrat, Vladlen Koltun, and Micha Sharir, "On the Union of κ -Round Objects in Three and Four Dimensions." *Discrete and Computational Geometry*, 36, 511–526, 2006. Special issue on the 20th ACM Symposium on Computational Geometry.
95. Pankaj K. Agarwal, Boris Aronov, and Vladlen Koltun, "Efficient Algorithms for Bichromatic Separability." *ACM Transactions on Algorithms*, 2, 209–227, 2006. Also in 15th ACM-SIAM Symposium on Discrete Algorithms.
96. Pankaj K. Agarwal, Boris Aronov, Vladlen Koltun, and Micha Sharir, "On Lines Avoiding Unit Balls in Three Dimensions." *Discrete and Computational Geometry*, 34, 231–250, 2005. Also in 20th Symposium on Computational Geometry.
97. Vladlen Koltun and Micha Sharir, "Curve-sensitive Cuttings," *SIAM Journal on Computing*, 34(4): 863–878, 2005. Also in 19th Symposium on Computational Geometry.
98. Boris Aronov, Vladlen Koltun, and Micha Sharir, "Cutting Triangular Cycles of Lines in Space," *Discrete and Computational Geometry*, 34(2): 231–250, 2005. Also in 35th ACM Symposium on Theory of Computing.
99. Boris Aronov, Vladlen Koltun, and Micha Sharir, "Incidences Between Points and Circles in Three Dimensions," *Discrete and Computational Geometry*, 33(2): 185–206, 2005. Also in 18th Symposium on Computational Geometry.

100. Vladlen Koltun and Carola Wenk, "Matching Polyhedral Terrains Using Overlays of Envelopes," *Algorithmica*, 41(3): 159–183, 2005.
101. Vladlen Koltun, "Almost Tight Upper Bounds for Vertical Decompositions in Four Dimensions," *Journal of the ACM*, 51, 699–730, 2004. Also in 42nd IEEE Symposium on Foundations of Computer Science.
102. Vladlen Koltun, "Sharp Bounds for Vertical Decompositions of Linear Arrangements in Four Dimensions," *Discrete and Computational Geometry*, 31, 435–460, 2004.
103. Vladlen Koltun and Micha Sharir, "Polyhedral Voronoi Diagrams of Polyhedra in Three Dimensions," *Discrete and Computational Geometry*, 31, 83–124. Special issue on the 18th ACM Symposium on Computational Geometry.
104. Vladlen Koltun, "Ready, Set, Go! - The Voronoi Diagram of Moving Points that Start from a Line," *Information Processing Letters*, 89, 233–235, 2004.
105. Vladlen Koltun, "Segment Intersection Searching Problems in General Settings," *Discrete and Computational Geometry*, 30, 25–44. Special issue on the 17th ACM Symposium on Computational Geometry.
106. Vladlen Koltun and Micha Sharir, "The Partition Technique for Overlays of Envelopes," *SIAM Journal on Computing*, 32, 841–863, 2003. Also in 43rd IEEE Symposium on Foundations of Computer Science.
107. Vladlen Koltun and Micha Sharir, "3-Dimensional Euclidean Voronoi Diagrams of Lines with a Fixed Number of Orientations," *SIAM Journal on Computing*, 32, 616–642, 2003. Also in 18th Symposium on Computational Geometry.

Technical reports (selected):

108. Peter Anderson, Angel Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, Amir R. Zamir, "On Evaluation of Embodied Navigation Agents," Technical Report, arXiv:1807.06757, 2018.
109. François Germain, Qifeng Chen, and Vladlen Koltun, "Speech Denoising with Deep Feature Losses," Technical Report, arXiv:1806.10522, 2018.
110. Shaojie Bai, J. Zico Kolter, and Vladlen Koltun, "An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling," Technical Report, arXiv:1803.01271, 2018.
111. Sohil Atul Shah and Vladlen Koltun, "Deep Continuous Clustering," Technical Report, arXiv:1803.01449, 2018.
112. Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun, "Open3D: A Modern Library for 3D Data Processing," Technical Report, arXiv:1801.09847, 2018.
113. Manolis Savva, Angel X. Chang, Alexey Dosovitskiy, Thomas Funkhouser, and Vladlen Koltun, "MINOS: Multimodal Indoor Simulator for Navigation in Complex Environments," Technical Report, arXiv:1712.03931, 2017.

114. Sungjoon Choi, Qian-Yi Zhou, Stephen Miller, and Vladlen Koltun, "A Large Dataset of Object Scans," Technical Report, arXiv:1602.02481, 2016.