

Vladlen Koltun
Curriculum Vitae

vladlen.info
vladlen.koltun@gmail.com
Born Nov 1980 in Kyiv, Ukraine. US citizen.

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

- | | |
|------------------|--|
| 8/2021 – present | Apple, USA
<i>Distinguished Scientist</i> <ul style="list-style-type: none">• Direct a new international research organization. |
| 10/2019 – 8/2021 | Intel, USA
<i>Chief Scientist for Intelligent Systems</i> <ul style="list-style-type: none">• Named Chief Scientist for Intelligent Systems at Intel.• Grew the Intelligent Systems Lab to ~30 people on four continents.• Continued leading machine learning, robotics, and computer vision research.<ul style="list-style-type: none">○ In machine learning, "Multiscale Deep Equilibrium Models" was selected for oral presentation at the Neural Information Processing Systems 2020 symposium (1.1% selection rate).○ In robotics, "Learning Quadrupedal Locomotion over Challenging Terrain" was published in Science Robotics in October 2020, featured on the cover, and covered widely in the scientific press as a significant milestone. "Deep Drone Acrobatics" was nominated for the Best Paper Award at the Robotics: Science and Systems 2020 conference, one of 3 papers nominated for the award out of 321 submissions.○ In computer vision, "Habitat: A Platform for Embodied AI Research" was nominated for the Best Paper Award at the International Conference on Computer Vision 2019, one of 11 papers nominated or selected for awards out of 4303 submissions.• Ramped up research in computational science, starting with applications of deep networks and differentiable programming to simulation of physical systems.• Led a task force of research leaders across industry and academia that formulated a program for research in Embodied AI over the next 5-10 years, presented in a widely circulated report.• Continued to develop widely used open-source projects.<ul style="list-style-type: none">○ Open3D was expanded significantly with new machine learning and visualization capabilities. The Open3D codebase was downloaded more |

than a million times in 2020.

- CARLA saw four new releases in 2020; the CARLA Leaderboard went live in summer 2020 and the second CARLA Challenge was organized at the Neural Information Processing Systems 2020 symposium. The CARLA codebase was downloaded more than two million times in H2 2020 alone.
- Launched OpenBot, a new open-source hardware and software initiative in robotics.

3/2018 – 9/2019 Intel, Santa Clara, CA

Senior Principal Researcher

Renamed the lab to Intelligent Systems Lab reflect its broad scope. The lab conducts high-impact basic research on intelligent systems, covering machine learning, robotics, computer vision, and graphics. Continued to grow the lab. Supervised two dozen interns in two years. Expanded to a third location (Bangalore, India). Hired research engineers to support larger projects in the lab. Significantly expanded the lab's computational infrastructure. Ramped up robotics research (Best Systems Paper Award at the Conference on Robot Learning; two papers in Science Robotics). Strengthened work on core machine learning in the lab. Led an industry-wide working group on empirical methodology in navigation research, resulting in widely adopted recommendations. Co-created Habitat, a platform for research in embodied AI. (Habitat was inspired by MINOS, an earlier simulation platform.) Continued to direct the Intel Network on Intelligent Systems; this is the highest-profile Intel investment in AI research in academia and brings together the leading academic labs in robotics, computer vision, and machine learning in Europe. Further ramped up research on deep learning for computational imaging, with high-impact results on low-light imaging. Continued progress on deep learning for 3D data processing. Continued investment in large-scale projects such as Open3D and CARLA.

1/2015 – 2/2018 Intel, Santa Clara, CA

Principal Researcher

Built a research lab (Visual Computing 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 <i>Research Fellow</i> Conducted research in theoretical computer science.
9/2002 – 6/2005	University of California, Berkeley, CA <i>Postdoctoral Researcher</i> Conducted research in theoretical computer science. Supervisor: Christos Papadimitriou
6/1997 – 5/1999	Shells Interactive, Israel <i>Senior Software Developer</i>

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 Morgenthaller 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

- "SPEAR: A Simulator for Photorealistic Embodied AI Research"
European Conference on Computer Vision (ECCV), 2026.
Selected for a spotlight oral presentation at the conference. Selection rate 1.6%. (164 papers selected for spotlight oral or full oral, out of 10,473 submissions.)

- *"Cut Your Losses in Large-Vocabulary Language Models"*
International Conference on Learning Representations (ICLR), 2025.
Selected for oral presentation at the conference. Selection rate 1.8%.
- *"Champion-level drone racing using deep reinforcement learning"*
Nature, 620, 2023.
Featured on the cover of the magazine.
- *"Reaching the limit in autonomous racing: Optimal control versus reinforcement learning"*
Science Robotics, 8(82), 2023.
Featured on the cover of the magazine.
- ***"Optimizing Locomotion Controllers Using Biologically-Based Actuators and Objectives"***
ACM SIGGRAPH Test-of-Time Award at SIGGRAPH 2023.
The award "recognizes highly influential papers published in SIGGRAPH conferences that have made a significant impact over the past 10 years or more". This is the first year of this annual award, for which all papers presented at SIGGRAPH conferences from 2011 to 2013 were considered. Four papers were selected for the award.
- *"Guaranteed Conservation of Momentum for Learning Particle-based Fluid Dynamics"*
Neural Information Processing Systems (NeurIPS), 2022.
Selected for oral presentation at the conference.
- *"Dancing under the stars: video denoising in starlight"*
Computer Vision and Pattern Recognition (CVPR), 2022.
Selected for oral presentation at the conference. Selection rate 4.2%. (344 papers selected out of 8161 submissions.)
- *"Habitat 2.0: Training Home Assistants to Rearrange their Habitat"*
Advances in Neural Information Processing Systems (NeurIPS), 2021.
Selected for spotlight at the conference. Selection rate <3%.
- *"Point Transformer"*
International Conference on Computer Vision (ICCV), 2021.
Selected for oral presentation at the conference. Selection rate 3.4%. (210 papers selected out of 6236 submissions.)
- *"Learning to Drive from a World on Rails"*
International Conference on Computer Vision (ICCV), 2021.
Selected for oral presentation at the conference. Selection rate 3.4%. (210 papers selected out of 6236 submissions.)
- *"Self-supervised Geometric Perception"*
Computer Vision and Pattern Recognition (CVPR), 2021.
Selected for oral presentation at the conference.

- ***“Deep Drone Racing: From Simulation to Reality With Domain Randomization”***
IEEE Transactions on Robotics, 2020.
Selected for Honorable Mention for the IEEE Transactions on Robotics Best Paper Award. 4 papers were selected, out of all papers published in the journal that year.
- ***“Multiscale Deep Equilibrium Models”***
Neural Information Processing Systems (NeurIPS), 2020.
Selected for oral presentation at the conference. Selection rate 1.1%. (105 papers selected out of 9454 submissions.)
- ***“Tracking Objects as Points”***
European Conference on Computer Vision (ECCV), 2020.
Selected for a spotlight oral presentation at the conference. Selection rate 5.3%. (265 papers selected for spotlight oral or full oral, out of 5025 submissions.)
- ***“Learning Quadrupedal Locomotion over Challenging Terrain”***
Science Robotics, 5(47), 2020.
Featured on the cover of the magazine.
- ***“Deep Drone Acrobatics”***
Robotics: Science and Systems (RSS), 2020.
Nominated for the Best Paper Award. 3 papers were nominated for the award, out of 321 submissions to the conference.
- ***“Deep Global Registration”***
Computer Vision and Pattern Recognition (CVPR), 2020.
Selected for oral presentation at the conference. Selection rate 5.7%. (335 papers selected out of 5865 submissions.)
- ***“High-dimensional Convolutional Networks for Geometric Pattern Recognition”***
Computer Vision and Pattern Recognition (CVPR), 2020.
Selected for oral presentation at the conference. Selection rate 5.7%. (335 papers selected out of 5865 submissions.)
- ***“Learning to Control PDEs with Differentiable Physics”***
International Conference on Learning Representations (ICLR), 2020.
Selected for a spotlight oral presentation at the conference. Selection rate 6%. (156 papers selected for spotlight oral or full oral, out of 2594 submissions.)
- ***“Deep Equilibrium Models”***
Neural Information Processing Systems (NeurIPS), 2019.
Selected for a spotlight oral presentation at the conference. Selection rate 3%. (200 papers selected for spotlight oral or full oral, out of 6743 submissions.)
- ***“Habitat: A Platform for Embodied AI Research”***
International Conference on Computer Vision (ICCV), 2019.
Selected for full oral presentation and **nominated for the Best Paper Award. 11 papers were nominated or selected for awards, out of 4303 submissions to the conference.**
- ***“Seeing Motion in the Dark”***
International Conference on Computer Vision (ICCV), 2019.
Selected for full oral presentation at the conference. Selection rate 4.6%. (200 papers selected out of 4303 submissions.)

- *"Consensus Maximization Tree Search Revisited"*
International Conference on Computer Vision (ICCV), 2019.
Selected for full oral presentation at the conference. Selection rate 4.6%. (200 papers selected out of 4303 submissions.)
- *"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.)

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 k -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.

- Distinguished Lecture, ETH Zürich. 2026.
- Distinguished Colloquium, Princeton University Computer Science. 2023.
- Keynote speaker, *CVPR Workshop: Autonomous Driving*. 2021.
- Keynote speaker, *Eurographics*. 2021.
- Keynote speaker, *German Conference on Pattern Recognition*. 2020.
- Invited speaker, *ICML Workshop: AI for Autonomous Driving*. 2020.
- Keynote speaker, *CVPR Workshop: DeepVision*. 2020.
- Invited speaker, *Machines Can See*. Moscow, Russia, 2020.
- Keynote speaker, *Autonomous Vehicles and Machines*. Burlingame, CA, 2020.
- Invited speaker, *NeurIPS Workshop: Machine Learning for Autonomous Driving*. Vancouver, Canada, 2019.
- Invited speaker, *NeurIPS Workshop: Learning Transferable Skills*. Vancouver, Canada, 2019.
- Invited speaker, *ICCV Workshop: Advances in Image Manipulation*. Seoul, Korea, 2019.
- Invited speaker, *ICCV Workshop: 3D Reconstruction in the Wild*. Seoul, Korea, 2019.
- Invited speaker, *ICCV Workshop: Geometry Meets Deep Learning*. Seoul, Korea, 2019.
- Invited speaker, *ICCV Workshop: Deep Learning for Visual SLAM*. Seoul, Korea, 2019.
- Invited speaker, *Image Sensors Americas*. San Jose, CA, 2019.
- Invited speaker, *Scenes from Video IV*. San Bernardo, Spain, 2019.
- 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 senior area chair roles, area chair roles, program committee memberships, etc.

- Associate Editor-in-Chief, *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2023 – 2025
- Program Chair, *Computer Vision and Pattern Recognition (CVPR)*, 2023
- Outstanding Paper Award Committee, *International Conference on Learning Representations (ICLR)*, 2023

- Executive Best Paper Award Committee, *International Conference on Learning Representations (ICLR)*, 2020
- Editorial board, *SIAM Journal on Imaging Sciences*, 2020 – 2022
- Editorial board, *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2016 – 2023
- 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

This section was only maintained up to 2021.

Research scientists

- Dr. Benjamin Ummenhofer. Intel Labs. 2021.
- Dr. Ozan Sener. Intel Labs. 2019 – 2021. (Next: Apple.)
- Dr. Gernot Riegler. Intel Labs. 2019 – 2020. (Next: Amazon.)
- Dr. Matthias Müller. Intel Labs. 2019 – 2021. (Next: Apple.)
- Dr. Sohil Shah. Intel Labs. 2018 – 2021. (Next: startup.)
- Dr. German Ros. Intel Labs. 2018 – 2021. (Next: Nvidia.)
- Dr. Stephan Richter. Intel Labs. 2018 – 2021. (Next: Apple.)
- Dr. Qifeng Chen. Intel Labs. 2017 – 2018. (Next: Faculty at HKUST.)
- Dr. Katelyn Gao. Intel Labs. 2017 – 2021.
- Dr. René Ranftl. Intel Labs. 2017 – 2021. (Next: Epic Games.)
- Dr. Alexey Dosovitskiy. Intel Labs. 2017 – 2019. (Next: Google Brain.)
- Dr. Philipp Krähenbühl. Intel Labs. 2016 – 2021. (Part time.)
- Dr. Jaesik Park. Intel Labs. 2015 – 2019. (Next: Faculty at POSTECH, then Seoul National University.)
- 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, then Microsoft Research.)
- Dr. Qian-Yi Zhou. Intel Labs. 2015 – 2017. (Next: Co-founded a start-up with experienced entrepreneurs, later sold to Snap.)

Postdoctoral researchers

- Dr. Titus Cieslewski. Intel Labs. 2021.

- Dr. Samarth Brahmbhatt. Intel Labs. 2020 – 2021.
- Dr. Zhipeng Cai. Intel Labs. 2020 – 2021.
- Dr. Benjamin Ummenhofer. Intel Labs. 2018 – 2020.
- Dr. Ozan Sener. Intel Labs. 2018 – 2019.
- Dr. Gernot Riegler. Intel Labs. 2018 – 2019.
- 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.)

- Nando Metzger, 2026, ETH Zürich. Supervisor: Konrad Schindler.
- Swaminathan Gurumurthy, 2025, CMU. Supervisor: Zico Kolter.
- Takahiro Miki, 2025, ETH Zürich. Supervisor: Marco Hutter.
- Naoki Yokoyama, 2025, Georgia Tech. Supervisors: Dhruv Batra and Sehoon Ha.
- Arjun Majumdar, 2024, Georgia Tech. Supervisor: Dhruv Batra.
- Joanne Truong, 2024, Georgia Tech. Supervisors: Dhruv Batra and Sonia Chernova.

- Joonho Lee, 2024, ETH Zürich. Supervisor: Marco Hutter.
- Wei Dong, 2023, CMU. Supervisor: Michael Kaess. (After PhD: Apple)
- Erik Wijmans, 2022, Georgia Tech. Supervisors: Dhruv Batra and Irfan Essa. (After PhD: Apple)
- Ankit Goyal, 2022, Princeton. Supervisor: Jia Deng. (After PhD: NVIDIA)
- Shaojie Bai, 2022, CMU. Supervisor: Zico Kolter. (After PhD: Meta)
- Marc Khoury, 2020, UC Berkeley. Supervisor: Jonathan Shewchuk. (After PhD: Hudson River Trading)
- Carl Toft, 2020, Chalmers University of Technology. Supervisor: Fredrik Kahl. (Faculty opponent.)
- Jemin Hwangbo, 2018, ETH Zürich. Supervisor: Marco Hutter. (After PhD: Faculty at KAIST)
- Brandon Amos, 2018, CMU. Supervisor: Zico Kolter. (After PhD: Facebook AI Research)
- Chen Chen, 2018, UIUC. Supervisor: Minh Do. (After PhD: Apple)
- 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, then faculty at ETH Zürich)

Doctoral interns

- Zachary Teed. Intel Labs, visiting from Princeton. 2021. (After PhD: Skydio.)
- Andrew Spielberg. Intel Labs, visiting from MIT. 2021. (After PhD: Postdoc at Harvard.)
- Ameya Prabhu. Intel Labs, visiting from University of Oxford. 2021.
- Zhaoming Xie. Intel Labs, visiting from UBC. 2021. (After PhD: Postdoc at Stanford.)
- Boyi Li. Intel Labs, visiting from Cornell. 2021. (After PhD: Research Scientist at NVIDIA and Postdoc at UC Berkeley)
- Shuyang Sun. Intel Labs, visiting from University of Oxford. 2021.
- Patrick Stotko. Intel Labs, visiting from University of Bonn. 2021. (After PhD: Postdoc at Bonn.)
- Guandao Yang. Intel Labs, visiting from Cornell. 2021. (After PhD: Postdoc at Stanford.)
- Ankit Goyal. Intel Labs, visiting from Princeton. 2021. (After PhD: Research Scientist at NVIDIA.)
- Feihu Zhang. Intel Labs, visiting from University of Oxford. 2021.
- Kristina Monakhova. Intel Labs, visiting from UC Berkeley. 2021. (After PhD: Postdoc at MIT, then Assistant Professor at Cornell.)
- Lukas Prantl. Intel Labs, visiting from TU München. Fall/winter 2020/21.
- Guohao Li. Intel Labs, visiting from KAUST. Fall/winter 2020/21.
- Pan Liu. Intel Labs, visiting from Penn State. Fall/winter 2020/21. (After PhD: Intel Labs.)
- Yiling Qiao. Intel Labs, visiting from University of Maryland. Summer 2020.
- Erik Wijmans. Intel Labs, visiting from Georgia Tech. Summer 2020. (After PhD: Research Scientist at Apple.)

- Heng Yang. Intel Labs, visiting from MIT. Summer 2020. (After PhD: Assistant Professor at Harvard.)
- Kai Zhang. Intel Labs, visiting from Cornell. Summer 2020. Second internship Summer 2021. (After PhD: Research Scientist at Adobe.)
- Johan Björck. Intel Labs, visiting from Cornell. Summer 2019. (After PhD: Microsoft.)
- Wei Dong. Intel Labs, visiting from CMU. Summer 2019. Second internship Summer 2020. (After PhD: Apple.)
- Hexiang (Frank) Hu. Intel Labs, visiting from USC. Summer 2019. (After PhD: Research Scientist at Google Brain.)
- John Lambert. Intel Labs, visiting from Georgia Tech. Summer 2019. (After PhD: Research Scientist at Waymo.)
- Zhuang Liu. Intel Labs, visiting from UC Berkeley. Summer 2019. (After PhD: Meta AI Research.)
- Aleksei Petrenko. Intel Labs, visiting from USC. Summer 2019. Second internship Summer 2020. (After PhD: Research Scientist at Apple.)
- Anirudh Vemula. Intel Labs, visiting from CMU. Summer 2019. (After PhD: Aurora Innovation.)
- Xingyi Zhou. Intel Labs, visiting from UT Austin. Summer 2019. Second internship Summer 2020. (After PhD: Google Research.)
- Christopher Choy. Intel Labs, visiting from Stanford. Winter/spring 2019. (After PhD: Research Scientist at NVIDIA.)
- Hassan Abu AlHaija. Intel Labs, visiting from University of Heidelberg. Winter/spring 2019. (After PhD: PCH Innovations.)
- Hengshuang Zhao. Intel Labs, visiting from CUHK. Winter/spring 2019. (After PhD: Assistant Professor, University of Hong Kong.)
- Zhipeng Cai. Intel Labs, visiting from University of Adelaide. Winter/spring 2019. (After PhD: Intel Labs.)
- Richard Shin. Intel Labs, visiting from UC Berkeley. Fall/winter 2018/2019. (After PhD: Microsoft Research.)
- Brandon Amos. Intel Labs, visiting from CMU. Summer/fall 2018. (After PhD: Facebook AI Research.)
- Shaojie Bai. Intel Labs, visiting from CMU. Summer/fall 2018.
- Adel Bibi. Intel Labs, visiting from KAUST. Summer/fall 2018. (After PhD: Postdoc, University of Oxford.)
- Yang Gao. Intel Labs, visiting from UC Berkeley. Summer/fall 2018. (After PhD: Assistant Professor, Tsinghua University.)
- Katrin Lasinger. Intel Labs, visiting from ETH Zürich. Summer/fall 2018. (After PhD: Facebook Reality Labs.)
- David Lindell. Intel Labs, visiting from Stanford. Summer/fall 2018. (After PhD: Assistant Professor, University of Toronto.)
- Dmytro Mishkin. Intel Labs, visiting from Czech Technical University. Summer/fall 2018. (After PhD: Postdoc, Czech Technical University.)
- Henri Rebecq. Intel Labs, visiting from University of Zürich. Summer/fall 2018. (After PhD: Research Scientist at Google.)
- Abhay Yadav. Intel Labs, visiting from University of Maryland. Summer/fall 2018.

- Xuaner Zhang. Intel Labs, visiting from UC Berkeley. Summer/fall 2018. (After PhD: Research Scientist at Adobe.)
- Matthias Vestner. Intel Labs, visiting from TU München. October 2017 – April 2018. (After PhD: Apple.)
- Nikolay Savinov. Intel Labs, visiting from ETH Zürich. June-December 2017. (After PhD: DeepMind.)
- Maxim Tatarchenko. Intel Labs, visiting from Freiburg. May-November 2017. (After PhD: Bosch Center for Artificial Intelligence.)
- Sergey Zagoruyko. Intel Labs, visiting from ParisTech. September 2017 – February 2018. (After PhD: Facebook AI Research.)
- Xiaojuan Qi. Intel Labs, visiting from CUHK. May-November 2017. (After PhD: Postdoc at the University of Oxford, then faculty at the University of Hong Kong.)
- François Germain. Intel Labs, visiting from Stanford. June 2017 – March 2018. (After PhD: Mitsubishi Electric Research Laboratories.)
- 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. (After PhD: Postdoc at MILA.)
- 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. (After PhD: CSIRO.)
- 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. (After PhD: Hudson River Trading)
- Mohammad Haris Baig. Intel Labs, visiting from Dartmouth. June-December 2015. (After PhD: Apple.)
- Linguang Zhang. Intel Labs, visiting from Princeton. June-September 2015. (After PhD: Facebook Reality Labs.)
- Jakob Engel. Intel Labs, visiting from TU München. May-November 2015. (After PhD: Facebook Reality Labs.)
- 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.)

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. Mike Roberts, Renhan Wang, Rushikesh Zavar, Rachith Dey-Prakash, Quentin Leboutet, Stephan R. Richter, Matthias Müller, German Ros, Rui Tang, Stefan Leutenegger, Yannick Hold-Geoffroy, Kalyan Sunkavalli, and Vladlen Koltun, “SPEAR: A Simulator for Photorealistic Embodied AI Research,” *European Conference on Computer Vision (ECCV)*, September 2026
2. Lars Mescheder, Wei Dong, Shiwei Li, Xuyang Bai, Marcel Santos, Peiyun Hu, Bruno Lecouat, Mingmin Zhen, Amaël Delaunoy, Tian Fang, Yanghai Tsin, Stephan R. Richter, and Vladlen Koltun, “Sharp Monocular View Synthesis in Less Than a Second,” *International Conference on Learning Representations (ICLR)*, April 2026
3. Marco Cusumano-Towner, David Hafner, Alex Hertzberg, Brody Huval, Aleksei Petrenko, Eugene Vinitsky, Erik Wijmans, Taylor Killian, Stuart Bowers, Ozan Sener, Philipp Krähenbühl, and Vladlen Koltun, “Robust Autonomy Emerges from Self-Play,” *International Conference on Machine Learning (ICML)*, July 2025
4. Santhosh Kumar Ramakrishnan, Erik Wijmans, Philipp Krähenbühl, and Vladlen Koltun, “Does Spatial Cognition Emerge in Frontier Models?,” *International Conference on Learning Representations (ICLR)*, April 2025
5. Aleksei Bochkovskii, Amaël Delaunoy, Hugo Germain, Marcel Santos, Yichao Zhou, Stephan R. Richter, and Vladlen Koltun, “Depth Pro: Sharp Monocular Metric Depth in Less Than a Second,” *International Conference on Learning Representations (ICLR)*, April 2025
6. Alejandro Newell, Peiyun Hu, Lahav Lipson, Stephan Richter, and Vladlen Koltun, “CoMotion: Concurrent Multi-person 3D Motion,” *International Conference on Learning Representations (ICLR)*, April 2025
7. Erik Wijmans, Brody Huval, Alexander Hertzberg, Vladlen Koltun, and Philipp Krähenbühl, “Cut Your Losses in Large-Vocabulary Language Models,” *International Conference on Learning Representations (ICLR)*, April 2025
8. Matthias Müller, Samarth Brahmabhatt, Ankur Deka, Quentin Leboutet, David Hafner, and Vladlen Koltun, “OpenBot-Fleet: A System for Collective Learning with Real Robots,” *International Conference on Robotics and Automation (ICRA)*, May 2024
9. Diana Wofk, René Ranftl, Matthias Müller, and Vladlen Koltun, “Monocular Visual-Inertial Depth Estimation,” *International Conference on Robotics and Automation (ICRA)*, May 2023
10. Ankit Goyal, Alexey Bochkovskiy, Jia Deng, and Vladlen Koltun, “Non-deep Networks,” *Advances in Neural Information Processing Systems (NeurIPS)*, December 2022
11. Ozan Sener and Vladlen Koltun, “Domain Generalization without Excess Empirical Risk,” *Advances in Neural Information Processing Systems (NeurIPS)*, December 2022

12. Lukas Prantl, Benjamin Ummenhofer, Vladlen Koltun, and Nils Thuerey, "Guaranteed Conservation of Momentum for Learning Particle-based Fluid Dynamics," *Advances in Neural Information Processing Systems (NeurIPS)*, December 2022
13. Philipp Holl, Vladlen Koltun, and Nils Thuerey, "Scale-invariant Learning by Physics Inversion," *Advances in Neural Information Processing Systems (NeurIPS)*, December 2022
14. Kristina Monakhova, Stephan Richter, Laura Waller, and Vladlen Koltun, "Dancing under the stars: video denoising in starlight", *Computer Vision and Pattern Recognition (CVPR)*, June 2022
15. Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl, "Simple multi-dataset detection", *Computer Vision and Pattern Recognition (CVPR)*, June 2022
16. Xingyi Zhou, Tianwei Yin, Vladlen Koltun, and Philipp Krähenbühl, "Global Tracking Transformers", *Computer Vision and Pattern Recognition (CVPR)*, June 2022
17. Chenyang Lei, Chenyang Qi, Jiaxin Xie, Na Fan, Vladlen Koltun, and Qifeng Chen, "Shape from Polarization for Complex Scenes in the Wild", *Computer Vision and Pattern Recognition (CVPR)*, June 2022
18. Boyi Li, Kilian Q. Weinberger, Serge Belongie, Vladlen Koltun, and René Ranftl, "Language-driven Semantic Segmentation," *International Conference on Learning Representations (ICLR)*, April 2022
19. Shaojie Bai, Vladlen Koltun, and J. Zico Kolter, "Neural Deep Equilibrium Solvers," *International Conference on Learning Representations (ICLR)*, April 2022
20. Guandao Yang, Serge Belongie, Bharath Hariharan, and Vladlen Koltun, "Geometry Processing with Neural Implicit Fields," *Advances in Neural Information Processing Systems (NeurIPS)*, December 2021
21. Yiling Qiao, Junbang Liang, Vladlen Koltun, and Ming C. Lin, "Differentiable Simulation of Soft Multi-body Systems," *Advances in Neural Information Processing Systems (NeurIPS)*, December 2021
22. Andrew Szot, Alexander Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John M Turner, Noah D Maestre, Mustafa Mukadam, Devendra Singh Chaplot, Oleksandr Maksymets, Aaron Gokaslan, Vladimír Vondruš, Franziska Meier, Wojciech Galuba, Angel X Chang, Zsolt Kira, Vladlen Koltun, Jitendra Malik, Manolis Savva, and Dhruv Batra, "Habitat 2.0: Training Home Assistants to Rearrange their Habitat," *Advances in Neural Information Processing Systems (NeurIPS)*, December 2021
23. René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun, "Vision Transformers for Dense Prediction", *International Conference on Computer Vision (ICCV)*, October 2021
24. Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip Torr, and Vladlen Koltun, "Point Transformer", *International Conference on Computer Vision (ICCV)*, October 2021
25. Benjamin Ummenhofer and Vladlen Koltun, "Adaptive Surface Reconstruction with Multiscale Convolutional Kernels", *International Conference on Computer Vision (ICCV)*, October 2021
26. Dian Chen, Vladlen Koltun, and Philipp Krähenbühl, "Learning to Drive from a World on Rails", *International Conference on Computer Vision (ICCV)*, October 2021
27. Zhipeng Cai, Ozan Sener, and Vladlen Koltun, "Online Continual Learning with Natural Distribution Shifts: An Empirical Study with Visual Data", *International Conference on Computer Vision (ICCV)*, October 2021

28. Aleksei Petrenko, Erik Wijmans, Brennan Shacklett, and Vladlen Koltun, "Megaverse: Simulating Embodied Agents at One Million Experiences per Second", *International Conference on Machine Learning (ICML)*, July 2021
29. Yiling Qiao, Junbang Liang, Vladlen Koltun, and Ming C. Lin, "Efficient Differentiable Simulation of Articulated Bodies", *International Conference on Machine Learning (ICML)*, July 2021
30. Shaojie Bai, Vladlen Koltun, and J. Zico Kolter, "Stabilizing Equilibrium Models by Jacobian Regularization", *International Conference on Machine Learning (ICML)*, July 2021
31. Guohao Li, Matthias Müller, Bernard Ghanem, and Vladlen Koltun, "Training Graph Neural Networks with 1000 Layers", *International Conference on Machine Learning (ICML)*, July 2021
32. Gernot Riegler and Vladlen Koltun, "Stable View Synthesis", *Computer Vision and Pattern Recognition (CVPR)*, June 2021
33. Heng Yang, Wei Dong, Luca Carlone, and Vladlen Koltun, "Self-supervised Geometric Perception", *Computer Vision and Pattern Recognition (CVPR)*, June 2021
34. Matthias Müller and Vladlen Koltun, "OpenBot: Turning Smartphones into Robots", *International Conference on Robotics and Automation (ICRA)*, May 2021
35. Brennan Shacklett, Erik Wijmans, Aleksei Petrenko, Manolis Savva, Dhruv Batra, Vladlen Koltun, and Kayvon Fatahalian, "Large Batch Simulation for Deep Reinforcement Learning", *International Conference on Learning Representations (ICLR)*, May 2021
36. Shaojie Bai, Vladlen Koltun, and J. Zico Kolter, "Multiscale Deep Equilibrium Models", *Advances in Neural Information Processing Systems (NeurIPS)*, December 2020.
37. Gernot Riegler and Vladlen Koltun, "Free View Synthesis", *European Conference on Computer Vision (ECCV)*, August 2020
38. Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl, "Tracking Objects as Points", *European Conference on Computer Vision (ECCV)*, August 2020
39. Yiheng Chi, Abhiram Gnanasambandam, Vladlen Koltun, and Stanley H. Chan, "Dynamic Low-light Imaging with Quanta Image Sensors", *European Conference on Computer Vision (ECCV)*, August 2020
40. Aleksei Petrenko, Zhehui Huang, Tushar Kumar, Gaurav Sukhatme, and Vladlen Koltun, "Sample Factory: Egocentric 3D Control from Pixels at 100000 FPS with Asynchronous Reinforcement Learning", *International Conference on Machine Learning (ICML)*, July 2020
41. Yiling Qiao, Junbang Liang, Vladlen Koltun, and Ming Lin, "Scalable Differentiable Physics for Learning and Control", *International Conference on Machine Learning (ICML)*, July 2020
42. Elia Kaufmann, Antonio Loquercio, René Ranftl, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza, "Deep Drone Acrobatics", *Robotics: Science and Systems (RSS)*, July 2020
43. Hengshuang Zhao, Jiaya Jia, and Vladlen Koltun, "Exploring Self-attention for Image Recognition", *Computer Vision and Pattern Recognition (CVPR)*, Seattle, WA, June 2020.

44. Christopher Choy, Wei Dong, and Vladlen Koltun, "Deep Global Registration," *Computer Vision and Pattern Recognition (CVPR)*, Seattle, WA, June 2020.
45. Christopher Choy, Junha Lee, René Ranftl, Jaesik Park, and Vladlen Koltun, "High-dimensional Convolutional Networks for Geometric Pattern Recognition," *Computer Vision and Pattern Recognition (CVPR)*, Seattle, WA, June 2020.
46. Carolin Schmitt, Simon Donné, Gernot Riegler, Vladlen Koltun, and Andreas Geiger, "On Joint Estimation of Pose, Geometry and svBRDF from a Handheld Scanner," *Computer Vision and Pattern Recognition (CVPR)*, Seattle, WA, June 2020.
47. Ozan Sener and Vladlen Koltun, "Learning to Guide Random Search," *International Conference on Learning Representations (ICLR)*, Addis Ababa, Ethiopia, April 2020.
48. Philipp Holl, Vladlen Koltun, and Nils Thürey, "Learning to Control PDEs with Differentiable Physics," *International Conference on Learning Representations (ICLR)*, Addis Ababa, Ethiopia, April 2020.
49. Benjamin Ummenhofer, Lukas Prantl, Nils Thürey, and Vladlen Koltun, "Lagrangian Fluid Simulation with Continuous Convolutions," *International Conference on Learning Representations (ICLR)*, Addis Ababa, Ethiopia, April 2020.
50. Shaojie Bai, J. Zico Kolter, and Vladlen Koltun, "Deep Equilibrium Models," *Advances in Neural Information Processing Systems (NeurIPS)*, Vancouver, Canada, December 2019.
51. Junbang Liang, Ming Lin, and Vladlen Koltun, "Differentiable Cloth Simulation for Inverse Problems," *Advances in Neural Information Processing Systems (NeurIPS)*, Vancouver, Canada, December 2019.
52. Dian Chen, Brady Zhou, Vladlen Koltun, and Philipp Krähenbühl, "Learning by Cheating," *Conference on Robot Learning (CoRL)*, Osaka, Japan, October 2019.
53. Chen Chen, Qifeng Chen, Minh Do, and Vladlen Koltun, "Seeing Motion in the Dark," *International Conference on Computer Vision (ICCV)*, Seoul, Korea, October 2019.
54. Hyukryul Yang, Hao Ouyang, Vladlen Koltun, and Qifeng Chen, "Hiding Video in Audio via Reversible Generative Models," *International Conference on Computer Vision (ICCV)*, Seoul, Korea, October 2019.
55. Christopher Choy, Jaesik Park, and Vladlen Koltun, "Fully Convolutional Geometric Features," *International Conference on Computer Vision (ICCV)*, Seoul, Korea, October 2019.
56. Zhipeng Cai, Tat-Jun Chin, and Vladlen Koltun, "Consensus Maximization Tree Search Revisited," *International Conference on Computer Vision (ICCV)*, Seoul, Korea, October 2019.
57. Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, Dhruv Batra, "Habitat: A Platform for Embodied AI Research," *International Conference on Computer Vision (ICCV)*, Seoul, Korea, October 2019.
58. François Germain, Qifeng Chen, and Vladlen Koltun, "Speech Denoising with Deep Feature Losses," *Interspeech*, Graz, Austria, September 2019.
59. David Lindell, Gordon Wetzstein, and Vladlen Koltun, "Acoustic Non-Line-of-Sight Imaging," *Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, June 2019.

60. 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.
61. 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.
62. 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.
63. 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.
64. Shaojie Bai, J. Zico Kolter, and Vladlen Koltun, "Trellis Networks for Sequence Modeling," *International Conference on Learning Representations (ICLR)*, New Orleans, Louisiana, May 2019.
65. 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.
66. Ozan Sener and Vladlen Koltun, "Multi-Task Learning as Multi-Objective Optimization," *Advances in Neural Information Processing Systems (NIPS)*, Montréal, Canada, December 2018.
67. 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.
68. 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.
69. 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.
70. René Ranftl and Vladlen Koltun, "Deep Fundamental Matrix Estimation," *European Conference on Computer Vision (ECCV)*, Munich, Germany, September 2018.
71. 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.
72. 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.
73. 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.
74. 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.

75. 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.
76. Nikolay Savinov, Alexey Dosovitskiy, and Vladlen Koltun, "Semi-parametric Topological Memory for Navigation," *International Conference on Learning Representations (ICLR)*, Vancouver, Canada, May 2018.
77. 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.
78. 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.
79. 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.
80. 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.
81. Qifeng Chen and Vladlen Koltun, "Photographic Image Synthesis with Cascaded Refinement Networks," *International Conference on Computer Vision (ICCV)*, Venice, Italy, October 2017.
82. Stephan R. Richter, Zeeshan Hayder, and Vladlen Koltun, "Playing for Benchmarks," *International Conference on Computer Vision (ICCV)*, Venice, Italy, October 2017.
83. Qifeng Chen, Jia Xu, and Vladlen Koltun, "Fast Image Processing with Fully-Convolutional Networks," *International Conference on Computer Vision (ICCV)*, Venice, Italy, October 2017.
84. Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun, "Colored Point Cloud Registration Revisited," *International Conference on Computer Vision (ICCV)*, Venice, Italy, October 2017.
85. Marc Khoury, Qian-Yi Zhou, and Vladlen Koltun, "Learning Compact Geometric Features," *International Conference on Computer Vision (ICCV)*, Venice, Italy, October 2017.
86. Fisher Yu, Vladlen Koltun, and Thomas Funkhouser, "Dilated Residual Networks," *Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, July 2017.
87. 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.
88. Alexey Dosovitskiy and Vladlen Koltun, "Learning to Act by Predicting the Future," *International Conference on Learning Representations (ICLR)*, Toulon, France, April 2017.
89. Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun, "Fast Global Registration," *European Conference on Computer Vision (ECCV)*, Amsterdam, The Netherlands, October 2016.
90. 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.

91. 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.
92. 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.
93. 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.
94. Fisher Yu and Vladlen Koltun, "Multi-Scale Context Aggregation by Dilated Convolutions," *International Conference on Learning Representations (ICLR)*, San Juan, PR, May 2016.
95. Qifeng Chen and Vladlen Koltun, "Robust Nonrigid Registration by Convex Optimization," *International Conference on Computer Vision (ICCV)*, Santiago, Chile, December 2015.
96. Philipp Krähenbühl and Vladlen Koltun, "Learning to Propose Objects," *Computer Vision and Pattern Recognition (CVPR)*, Boston, MA, June 2015.
97. Sungjoon Choi, Qian-Yi Zhou, and Vladlen Koltun, "Robust Reconstruction of Indoor Scenes," *Computer Vision and Pattern Recognition (CVPR)*, Boston, MA, June 2015.
98. Qian-Yi Zhou and Vladlen Koltun, "Depth Camera Tracking with Contour Cues," *Computer Vision and Pattern Recognition (CVPR)*, Boston, MA, June 2015.
99. Philipp Krähenbühl and Vladlen Koltun, "Geodesic Object Proposals," *European Conference on Computer Vision (ECCV)*, Zürich, Switzerland, September 2014.
100. Sergey Levine and Vladlen Koltun, "Learning Complex Neural Network Policies with Trajectory Optimization," *International Conference on Machine Learning (ICML)*, Beijing, China, June 2014.
101. Qifeng Chen and Vladlen Koltun, "Fast MRF Optimization with Application to Depth Reconstruction," *Computer Vision and Pattern Recognition (CVPR)*, Columbus, OH, June 2014.
102. 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.
103. Sergey Levine and Vladlen Koltun, "Variational Guided Policy Search," *Advances in Neural Information Processing Systems (NIPS)*, Lake Tahoe, NV, December 2013.
104. 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.
105. Qian-Yi Zhou, Stephen Miller, and Vladlen Koltun, "Elastic Fragments for Dense Scene Reconstruction," *International Conference on Computer Vision (ICCV)*, Sydney, Australia, December 2013.
106. Philipp Krähenbühl and Vladlen Koltun, "Parameter Learning and Convergent Inference for Dense Random Fields," *International Conference on Machine Learning (ICML)*, Atlanta, GA, June 2013.
107. Sergey Levine and Vladlen Koltun, "Guided Policy Search," *International Conference on Machine Learning (ICML)*, Atlanta, GA, June 2013.
108. Philipp Krähenbühl and Vladlen Koltun, "Efficient Nonlocal Regularization for Optical Flow," *European*

Conference on Computer Vision (ECCV), Florence, Italy, October 2012.

109. Sergey Levine and Vladlen Koltun, "Continuous Inverse Optimal Control with Locally Optimal Examples," *International Conference on Machine Learning (ICML)*, Edinburgh, Scotland, June 2012.
110. 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.
111. 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.
112. 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.
113. Sergey Levine, Zoran Popovic and Vladlen Koltun, "Feature Construction for Inverse Reinforcement Learning," *Advances in Neural Information Processing Systems (NIPS)*, Vancouver, Canada, December 2010.
114. Sarel Har-Peled and Vladlen Koltun, "Separability with Outliers," *16th International Symposium on Algorithms and Computation*, Sanya, China, December 2005.
115. Jason Hartline and Vladlen Koltun, "Near-Optimal Pricing in Near-Linear Time," *9th Workshop on Algorithms and Data Structures*, Waterloo, Canada, August 2005.
116. Vladlen Koltun, "Pianos are not Flat: Rigid Motion Planning in Three Dimensions," *16th ACM-SIAM Symposium on Discrete Algorithms*, Vancouver, BC, January 2005.
117. 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.
118. Sarel Har-Peled, Vladlen Koltun, Dezheng Song, and Ken Goldberg, "Efficient Algorithms for Shared Camera Control," *19th ACM Symposium on Computational Geometry*, San Diego, CA, June 2003.
119. 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.
120. 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.
121. 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:

122. Yunlong Song, Angel Romero, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza, "Reaching the limit in autonomous racing: Optimal control versus reinforcement learning," *Science Robotics*, 8(82), 2023
123. Elia Kaufmann, Leonard Bauersfeld, Antonio Loquercio, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza, "Champion-level drone racing using deep reinforcement learning," *Nature*, 620, 2023
124. Brennan Shacklett, Luc Guy Rosenzweig, Zhiqiang Xie, Bidipta Sarkar, Andrew Szot, Erik Wijmans, Vladlen Koltun, Dhruv Batra, and Kayvon Fatahalian, "An Extensible, Data-Oriented Architecture for High-Performance, Many-World Simulation," *ACM Transactions on Graphics*, 42(4), 2023. Proceedings of ACM SIGGRAPH.
125. Wei Dong, Yixing Lao, Michael Kaess, and Vladlen Koltun, "ASH: A Modern Framework for Parallel Spatial Hashing in 3D Perception," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5), 2023.
126. Hexiang Hu, Ozan Sener, Fei Sha, and Vladlen Koltun, "Drinking from a Firehose: Continual Learning with Web-scale Natural Language," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5), 2023.
127. Stephan Richter, Hassan Abu AlHaija, and Vladlen Koltun, "Enhancing Photorealism Enhancement," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2), 2023.
128. John Lambert, Zhuang Liu, Ozan Sener, James Hays, and Vladlen Koltun, "MSeg: A Composite Dataset for Multi-domain Semantic Segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1), 2023. An earlier version appeared in Computer Vision and Pattern Recognition (CVPR) in 2020.
129. Takahiro Miki, Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Science Robotics*, 7(62), 2022.
130. Antonio Loquercio, Elia Kaufmann, René Ranftl, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza, "Learning high-speed flight in the wild," *Science Robotics*, 6(59), 2021.
131. Vladlen Koltun and David Hafner, "The h-index is no longer an effective correlate of scientific reputation," *PLOS ONE*, 16(6), 2021.
132. Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter, "Learning Quadrupedal Locomotion over Challenging Terrain," *Science Robotics*, 5(47), 2020.
133. René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun, "Towards Robust Monocular Depth Estimation: Mixing Datasets for Zero-Shot Cross-Dataset Transfer," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(3), 2022.
134. Henri Rebecq, René Ranftl, Vladlen Koltun, and Davide Scaramuzza, "High Speed and High Dynamic Range Video with an Event Camera," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(6), 2021. An earlier version appeared as "Events-to-Video: Bringing Modern Computer Vision to Event Cameras" in Computer Vision and Pattern Recognition (CVPR) in 2019.
135. Antonio Loquercio, Elia Kaufmann, René Ranftl, Alexey Dosovitskiy, Vladlen Koltun, and Davide Scaramuzza, "Deep Drone Racing: From Simulation to Reality With Domain Randomization," *IEEE Transactions on Robotics*, 36(1), 2020. An earlier version appeared as "Deep Drone Racing: Learning Agile

Flight in Dynamic Environments” in the Conference on Robot Learning (CoRL) in 2018.

136. Jan Carius, René Ranftl, Vladlen Koltun, and Marco Hutter, “Trajectory Optimization for Legged Robots with Slipping Motions,” *IEEE Robotics and Automation Letters*, 4(3), 2019.
137. Brady Zhou, Philipp Krähenbühl, and Vladlen Koltun, “Does Computer Vision Matter for Action?,” *Science Robotics*, 4(30), 2019.
138. 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.
139. 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.
140. Jan Carius, René Ranftl, Vladlen Koltun, and Marco Hutter, “Trajectory Optimization with Implicit Hard Contacts,” *IEEE Robotics and Automation Letters*, 3(4), 2018.
141. Jakob Engel, Vladlen Koltun, and Daniel Cremers, “Direct Sparse Odometry,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(3), 2018.
142. Sohil Atul Shah and Vladlen Koltun, “Robust Continuous Clustering,” *Proceedings of the National Academy of Sciences (PNAS)*, 114(37), 2017.
143. 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.
144. 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.
145. 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.
146. 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.
147. Qian-Yi Zhou and Vladlen Koltun, “Dense Scene Reconstruction with Points of Interest,” *ACM Transactions on Graphics*, 32(4), 2013. Proceedings of ACM SIGGRAPH.
148. 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.
149. 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.
150. 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.

151. 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.
152. 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.
153. 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.
154. 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.
155. 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.
156. 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.
157. 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.
158. 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.
159. Paul Merrell, Eric Schkufza, and Vladlen Koltun, "Computer-Generated Residential Building Layouts," *ACM Transactions on Graphics*, 29(6), 2010. Proceedings of ACM SIGGRAPH Asia.
160. Sergey Levine, Philipp Krähenbühl, Sebastian Thrun, and Vladlen Koltun, "Gesture Controllers," *ACM Transactions on Graphics*, 29(4), 2010. Proceedings of ACM SIGGRAPH.
161. 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.
162. 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.
163. Siddhartha Chaudhuri and Vladlen Koltun, "Smoothed Analysis of Probabilistic Roadmaps," *Computational Geometry*, 42(8): 731–747, 2009.
164. 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.
165. 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.

166. 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.
167. 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.
168. 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.
169. Vladlen Koltun and Micha Sharir, "Curve-sensitive Cuttings," *SIAM Journal on Computing*, 34(4): 863–878, 2005. Also in 19th Symposium on Computational Geometry.
170. 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.
171. 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.
172. Vladlen Koltun and Carola Wenk, "Matching Polyhedral Terrains Using Overlays of Envelopes," *Algorithmica*, 41(3): 159–183, 2005.
173. 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.
174. Vladlen Koltun, "Sharp Bounds for Vertical Decompositions of Linear Arrangements in Four Dimensions," *Discrete and Computational Geometry*, 31, 435–460, 2004.
175. 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.
176. Vladlen Koltun, "Ready, Set, Go! - The Voronoi Diagram of Moving Points that Start from a Line," *Information Processing Letters*, 89, 233–235, 2004.
177. 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.
178. 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.
179. 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):

180. Kevin Chen, Marco Cusumano-Towner, Brody Huval, Aleksei Petrenko, Jackson Hamburger, Vladlen Koltun, and Philipp Krähenbühl, “Reinforcement Learning for Long-Horizon Interactive LLM Agents,” Technical Report, arXiv:2502.01600, 2025
181. Feihu Zhang, Vladlen Koltun, Philip Torr, René Ranftl, and Stephan R Richter, “Unsupervised Contrastive Domain Adaptation for Semantic Segmentation,” Technical Report, arXiv:2204.08399, 2022
182. Kai Zhang, Gernot Riegler, Noah Snaveley, and Vladlen Koltun, “NeRF++: Analyzing and Improving Neural Radiance Fields,” Technical Report, arXiv:2010.07492, 2020.
183. Dhruv Batra, Angel X. Chang, Sonia Chernova, Andrew J. Davison, Jia Deng, Vladlen Koltun, Sergey Levine, Jitendra Malik, Igor Mordatch, Roozbeh Mottaghi, Manolis Savva, Hao Su, “Rearrangement: A Challenge for Embodied AI,” Technical Report, arXiv:2011.01975, 2020.
184. Artemij Amiranashvili, Nicolai Dorka, Wolfram Burgard, Vladlen Koltun, and Thomas Brox, “Scaling Imitation Learning in Minecraft”, Technical Report, arXiv:2007.02701, 2020.
185. Brandon Amos, Vladlen Koltun, and J. Zico Kolter, “The Limited Multi-Label Projection Layer,” Technical Report, arXiv:1906.08707, 2019.
186. Dmytro Mishkin, Alexey Dosovitskiy, and Vladlen Koltun, “Benchmarking Classic and Learned Navigation in Complex 3D Environments,” Technical Report, arXiv:1901.10915, 2019.
187. Charles Packer, Katelyn Gao, Jernej Kos, Philipp Krähenbühl, Vladlen Koltun, Dawn Song, “Assessing Generalization in Deep Reinforcement Learning,” Technical Report, arXiv:1810.12282, 2019
188. 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.
189. 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.
190. Sohil Atul Shah and Vladlen Koltun, “Deep Continuous Clustering,” Technical Report, arXiv:1803.01449, 2018.
191. Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun, “Open3D: A Modern Library for 3D Data Processing,” Technical Report, arXiv:1801.09847, 2018.
192. 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.
193. Sungjoon Choi, Qian-Yi Zhou, Stephen Miller, and Vladlen Koltun, “A Large Dataset of Object Scans,” Technical Report, arXiv:1602.02481, 2016.