

EDUCATION

- 2009–2014 **PhD, Computer Science**, *Stanford University*, Stanford CA
- 2007–2008 **Student Exchange, Computer Science**, *University of Washington*, Seattle WA
- 2005–2009 **BSc, Computer Science**, *ETH Zurich*, Zurich, Switzerland

RESEARCH AND INDUSTRY EXPERIENCE

- 2016– **Assistant Professor, Zeitlin Fellow**, *UT Austin*, Department of Computer Science
Computer Vision and Machine Learning
- 2014–2016 **Postdoctoral Researcher**, *UC Berkeley*, with Alexei Efros and Jitendra Malik
Computer Vision and Machine Learning
- 2009–2014 **Research Assistant**, *Stanford University*, with Vladlen Koltun
Computer Vision and Machine Learning
- 2012, 2013 **Research Internship**, *Adobe Creative Technologies Lab*
- 2011 **Research Internship**, *Disney Research Zurich*
- 2008 **Research Assistant**, *University of Washington*, with Zoran Popovic
- 2006–2009 **Research Assistant**, *ETH Zurich*, with Markus Gross, Gaston Gonnet and Bertrand Meyer

SERVICE AND AFFILIATIONS

- Area Chair International Conference on Computer Vision (ICCV): 2017
Computer Vision and Pattern Recognition (CVPR): 2018, 2019, 2020, 2021
European Conference on Computer Vision (ECCV): 2016, 2018
International Conference on Learning Representations (ICLR): 2021
Conference on Neural Information Processing Systems (NeurIPS): 2021

COLLABORATORS & OTHER AFFILIATIONS

- PhD Advisor Koltun, Vladlen; Stanford
- Postdoctoral Darrell, Trevor; UC-Berkeley
- Sponsors Efros, Alexei; UC-Berkeley
Malik, Jitendra; UC-Berkeley

HONORS AND AWARDS

- 2011 **NIPS Outstanding Student Paper Award**
Highest award at the conference, given to 3 out of 1400 submissions.
- 2010-2014 **Stanford Graduate Fellowship**
Pierre and Christine Lamond Fellow

GRANTS

- 2020-2025 **NSF AI Institute for Foundations of Machine Learning**
\$20,000,000; Senior Member (CO-PI)
- 2020-2023 **NSF SMALL: Towards inferring all properties of all things through time**
\$420,000; PI
- 2019-2024 **NSF CAREER: A Compression-Based Approach to Learning Video Representations**
\$500,000; PI
- 2016- **Industrial gifts: Intel, Berkeley Deep Drive**
\$338,000; PI

PUBLICATIONS

- [1] Dian Chen and Philipp Krähenbühl. Learning from all vehicles. In *CVPR*, 2022.
- [2] Brady Zhou and Philipp Krähenbühl. Cross-view transformers for real-time map-view semantic segmentation. In *CVPR*, 2022.
- [3] Xingyi Zhou, Tianwei Yin, Vladlen Koltun, and Philipp Krähenbühl. Global tracking transformers. In *CVPR*, 2022.
- [4] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Simple multi-dataset detection. In *CVPR*, 2022.
- [5] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Multimodal virtual point 3d detection. In *NeurIPS*, 2021.
- [6] Dian Chen, Vladlen Koltun, and Philipp Krähenbühl. Learning to drive from a world on rails. In *ICCV*, 2021.
- [7] Chao-Yuan Wu and Philipp Krähenbühl. Towards long-form video understanding. In *CVPR*, 2021.
- [8] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Center-based 3d object detection and tracking. In *CVPR*, 2021.
- [9] Aashaka Shah, Chao-Yuan Wu, Jayashree Mohan, Vijay Chidambaram, and Philipp Krähenbühl. Memory optimization for deep networks. In *ICLR*, 2021.
- [10] Brady Zhou, Nimit Kalra, and Philipp Krähenbühl. Domain adaptation through task distillation. In *ECCV*, 2020.
- [11] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking objects as points. In *ECCV*, 2020.
- [12] Chao-Yuan Wu, Ross Girshick, Kaiming He, Christoph Feichtenhofer, and Philipp Krähenbühl. A multigrid method for efficiently training video models. In *CVPR*, 2020.
- [13] Dian Chen, Brady Zhou, Vladlen Koltun, and Philipp Krähenbühl. Learning by cheating. In *CORL*, 2019.
- [14] Dequan Wang, Coline Devin, Qi-Zhi Cai, Philipp Krähenbühl, and Trevor Darrell. Monocular plan view networks for autonomous driving. In *IROS*, 2019.

- [15] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. In *arXiv preprint arXiv:1904.07850*, 2019.
- [16] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krähenbühl, and Ross Girshick. Long-term feature banks for detailed video understanding. In *CVPR*, 2019.
- [17] Xingyi Zhou, Jiacheng Zhuo, and Philipp Krähenbühl. Bottom-up object detection by grouping extreme and center points. In *CVPR*, 2019.
- [18] Hou-Ning Hu, Qi-Zhi Cai, Dequan Wang, Ji Lin, Min Sun, Philipp Krähenbühl, Trevor Darrell, and Fisher Yu. Joint monocular 3d vehicle detection and tracking. In *ICCV*, 2019.
- [19] Brady Zhou, Philipp Krähenbühl, and Vladlen Koltun. Does computer vision matter for action? In *Science Robotics*, 2019.
- [20] Brady Zhou and Philipp Krähenbühl. Don't let your discriminator be fooled. In *ICLR*, 2019.
- [21] Chao-Yuan Wu, Nayan Singhal, and Philipp Krähenbühl. Video compression through image interpolation. In *ECCV*, 2018. .
- [22] Haoshuo Huang, Qixing Huang, and Philipp Krähenbühl. Domain transfer through deep activation matching. In *ECCV*, 2018. .
- [23] Chao-Yuan Wu, Manzil Zaheer, Hexiang Hu, R. Manmatha, Alexander J. Smola, and Philipp Krähenbühl. Compressed video action recognition. In *CVPR*, 2018. .
- [24] Philipp Krähenbühl. Free supervision from video games. In *CVPR*, 2018. .
- [25] Chao-Yuan Wu, R. Manmatha, Alexander J. Smola, and Philipp Krähenbühl. Sampling matters in deep embedding learning. In *ICCV*, 2017. .
- [26] Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. In *ICLR*, 2017. .
- [27] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A. Efros. Generative visual manipulation on the natural image manifold. In *ECCV*, 2016. .
- [28] Deepak Pathak, Philipp Krähenbühl, Jeff Donahue, Trevor Darrell, and Alyosha Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016. .
- [29] Tinghui Zhou, Philipp Krähenbühl, Mathieu Aubry, Qixing Huang, and Alyosha Efros. Learning dense correspondence via 3d-guided cycle consistency. In *CVPR*, 2016. .
- [30] Philipp Krähenbühl, Carl Doersch, Jeff Donahue, and Trevor Darrell. Data-dependent initializations of convolutional neural networks. In *ICLR*, 2016.
- [31] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alyosha Efros. Learning a discriminative model for the perception of realism in composite images. In *ICCV*, 2015. .
- [32] Tinghui Zhou, Philipp Krähenbühl, and Alyosha Efros. Learning data-driven reflectance priors for intrinsic image decomposition. In *ICCV*, 2015. .
- [33] Deepak Pathak, Philipp Krähenbühl, and Trevor Darrell. Constrained convolutional neural networks for weakly supervised segmentation. In *ICCV*, 2015. .

- [34] Philipp Krähenbühl and Vladlen Koltun. Learning to propose objects. In *CVPR*, 2015. .
- [35] Philipp Krähenbühl and Vladlen Koltun. Geodesic object proposals. In *ECCV*, 2014. .
- [36] Philipp Krähenbühl and Vladlen Koltun. Parameter learning and convergent inference for dense random fields. In *ICML*, 2013.
- [37] Philipp Krähenbühl and Vladlen Koltun. Efficient nonlocal regularization for optical flow. In *ECCV*, 2012. .
- [38] Federico Perazzi, Philipp Krähenbühl, Yael Pritch, and Alexander Hornung. Saliency filters: Contrast based filtering for salient region detection. In *CVPR*, 2012. .
- [39] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NIPS*, 2011.
- [40] Sergey Levine, Philipp Krähenbühl, Sebastian Thrun, and Vladlen Koltun. Gesture controllers. In *SIGGRAPH*, 2010. .
- [41] Philipp Krähenbühl, Manuel Lang, Alexander Hornung, and Markus Gross. A system for retargeting of streaming video. In *SIGGRAPH Asia*, 2009. .

Complete list available at www.philkr.net.

Last CV update: June, 2022

GDC 4.824, 2317 Speedway – Austin, TX 78712 – USA

✉ [regular] philkr@cs.utexas.edu • 🌐 www.philkr.net • 🎓 [Google Scholar](#)

4/3