Philipp Krähenbühl

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

2022– **Associate Professor**, *UT Austin*, Department of Computer Science
2022– **Researcher**, *Apple*, AIML
2016–2022 **Assistant Professor**, *UT Austin*, Department of Computer Science
2014–2016 **Postdoctoral Researcher**, *UC Berkeley*, with Alexei Efros, Jitendra Malik, and Trevor Darrell
2016–2021 **Researcher (part-time)**, *Intel*
2009–2014 **Research Assistant**, *Stanford University*, with Vladlen Koltun
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 Sponsors
Darrell, Trevor; UC-Berkeley
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

GDC 4.824, 2317 Speedway – Austin, TX 78712 – USA
philkr@cs.utexas.edu • *www.philkr.net* • *Google Scholar*

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

2019-2026 **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] Shuhan Tan, Boris Ivanovic, Yuxiao Chen, Boyi Li, Xinshuo Weng, Yulong Cao, Philipp Krähenbühl, and Marco Pavone. Promptable closed-loop traffic simulation. In *CoRL*, 2024.

[2] Yue Zhao, Long Zhao, Xingyi Zhou, Jialin Wu, Chun-Te Chu, Hui Miao, Florian Schroff, Hartwig Adam, Ting Liu, Boqing Gong, Philipp Krähenbühl, and Liangzhe Yuan. Distilling vision-language models on millions of videos. In *CVPR*, 2024.

[3] Shuhan Tan, Boris Ivanovic, Xinshuo Weng, Marco Pavone, and Philipp Kraehenbuehl. Language conditioned traffic generation. In *CoRL*, 2023.

[4] Jeffrey Ouyang-Zhang, Daniel J. Diaz, Adam R. Klivans, and Philipp Krähenbühl. Predicting a protein's stability under a million mutations. In *NeurIPS*, 2023.

[5] Yue Zhao, Ishan Misra, Philipp Krähenbühl, and Rohit Girdhar. Learning video representations from large language models. In *CVPR*, 2023.

[6] Jang Hyun Cho, Philipp Krähenbühl, and Vignesh Ramanathan. Partdistillation: Learning parts from instance segmentation. In *CVPR*, 2023.

[7] Jang Hyun Cho and Philipp Krähenbühl. Language-conditioned detection transformer. In *CVPR*, 2023.

[8] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Simple multi-dataset detection. In *CVPR*, 2022.

[9] Yue Zhao and Philipp Krähenbühl. Real-time online video detection with temporal smoothing transformers. In *ECCV*, 2022.

[10] Brady Zhou and Philipp Krähenbühl. Cross-view transformers for real-time map-view semantic segmentation. In *CVPR*, 2022.

[11] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Philipp Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. In *ECCV*, 2022.

[12] Jang Hyun Cho and Philipp Krähenbühl. Long-tail detection with effective class-margins. In *ECCV*, 2022.

[13] Xingyi Zhou, Tianwei Yin, Vladlen Koltun, and Philipp Krähenbühl. Global tracking transformers. In *CVPR*, 2022.

[14] Dian Chen and Philipp Krähenbühl. Learning from all vehicles. In *CVPR*, 2022.

[15] Chao-Yuan Wu and Philipp Krähenbühl. Towards long-form video understanding. In *CVPR*, 2021.

[16] Aashaka Shah, Chao-Yuan Wu, Jayashree Mohan, Vijay Chidambaram, and Philipp Krähenbühl. Memory optimization for deep networks. In *ICLR*, 2021.

[17] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Center-based 3d object detection and tracking. In *CVPR*, 2021.

[18] Dian Chen, Vladlen Koltun, and Philipp Krähenbühl. Learning to drive from a world on rails. In *ICCV*, 2021.

[19] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Multimodal virtual point 3d detection. In *NeurIPS*, 2021.

[20] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Probabilistic two-stage detection. In *arXiv*, 2021.

[21] Brady Zhou, Nimit Kalra, and Philipp Krähenbühl. Domain adaptation through task distillation. In *ECCV*, 2020.

[22] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking objects as points. In *ECCV*, 2020.

[23] Chao-Yuan Wu, Ross Girshick, Kaiming He, Christoph Feichtenhofer, and Philipp Krähenbühl. A multigrid method for efficiently training video models. In *CVPR*, 2020.

[24] Dian Chen, Brady Zhou, Vladlen Koltun, and Philipp Krähenbühl. Learning by cheating. In *CORL*, 2019.

[25] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. In *arXiv preprint arXiv:1904.07850*, 2019.

[26] Brady Zhou, Philipp Krähenbühl, and Vladlen Koltun. Does computer vision matter for action? In *Science Robotics*, 2019.

[27] Brady Zhou and Philipp Krähenbühl. Don't let your discriminator be fooled. In *ICLR*, 2019.

[28] 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.

[29] 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.

[30] Dequan Wang, Coline Devin, Qi-Zhi Cai, Philipp Krähenbühl, and Trevor Darrell. Monocular plan view networks for autonomous driving. In *IROS*, 2019.

[31] Xingyi Zhou, Jiacheng Zhuo, and Philipp Krähenbühl. Bottom-up object detection by grouping extreme and center points. In *CVPR*, 2019.

[32] Chao-Yuan Wu, Manzil Zaheer, Hexiang Hu, R. Manmatha, Alexander J. Smola, and Philipp Krähenbühl. Compressed video action recognition. In *CVPR*, 2018.

[33] Chao-Yuan Wu, Nayan Singhal, and Philipp Krähenbühl. Video compression through image interpolation. In *ECCV*, 2018.

[34] Charles Packer, Katelyn Gao, Jernej Kos, Philipp Krähenbühl, Vladlen Koltun, and Dawn Song. Assessing generalization in deep reinforcement learning. In *arXiv*, 2018.

[35] Haoshuo Huang, Qixing Huang, and Philipp Krähenbühl. Domain transfer through deep activation matching. In *ECCV*, 2018.

[36] Philipp Krähenbühl. Free supervision from video games. In *CVPR*, 2018.

[37] Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. In *ICLR*, 2017.

[38] Chao-Yuan Wu, R. Manmatha, Alexander J. Smola, and Philipp Krähenbühl. Sampling matters in deep embedding learning. In *ICCV*, 2017.

[39] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A. Efros. Generative visual manipulation on the natural image manifold. In *ECCV*, 2016.

[40] Philipp Krähenbühl, Carl Doersch, Jeff Donahue, and Trevor Darrell. Data-dependent initializations of convolutional neural networks. In *ICLR*, 2016.

[41] Tinghui Zhou, Philipp Krähenbühl, Mathieu Aubry, Qixing Huang, and Alyosha Efros. Learning dense correspondence via 3d-guided cycle consistency. In *CVPR*, 2016.

[42] Deepak Pathak, Philipp Krähenbühl, Jeff Donahue, Trevor Darrell, and Alyosha Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016.

[43] 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.

[44] Philipp Krähenbühl and Vladlen Koltun. Learning to propose objects. In *CVPR*, 2015.

[45] Deepak Pathak, Philipp Krähenbühl, and Trevor Darrell. Constrained convolutional neural networks for weakly supervised segmentation. In *ICCV*, 2015.

[46] Tinghui Zhou, Philipp Krähenbühl, and Alyosha Efros. Learning data-driven reflectance priors for intrinsic image decomposition. In *ICCV*, 2015.

[47] Philipp Krähenbühl and Vladlen Koltun. Geodesic object proposals. In *ECCV*, 2014.

[48] Philipp Krähenbühl. Dense random fields. In *Phd Thesis*, 2014.

[49] Philipp Krähenbühl and Vladlen Koltun. Parameter learning and convergent inference for dense random fields. In *ICML*, 2013.

[50] Federico Perazzi, Philipp Krähenbühl, Yael Pritch, and Alexander Hornung. Saliency filters: Contrast based filtering for salient region detection. In *CVPR*, 2012.

[51] Philipp Krähenbühl and Vladlen Koltun. Efficient nonlocal regularization for optical flow. In *ECCV*, 2012.

[52] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NIPS*, 2011.

[53] Sergey Levine, Philipp Krähenbühl, Sebastian Thrun, and Vladlen Koltun. Gesture controllers. In *SIGGRAPH*, 2010.

[54] Philipp Krähenbühl. Art directable retargeting for streaming video. In *Bachelor Thesis, ETH Zurich*, 2009.

[55] 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: September, 2024