

RESEARCH INTERESTS Computer vision and machine learning, in particular deep learning for visual recognition and optimization during inference for robust modeling and adaptive computation.

I like to preserve classic ideas then make them differentiable. In my experience, reproducible research done right is extensible research, so I develop tools for do-it-yourself deep learning.

EDUCATION **University of California, Berkeley** · *Fall 2012 - Spring 2019*

Ph.D Computer Science

Advisor: Trevor Darrell

Thesis: Local and Adaptive Image-to-Image Learning and Inference

University of Massachusetts Amherst (Commonwealth College) · *Fall 2008 - Spring 2012*

B.S. Computer Science (AI concentration) & B.S. Psychology

with Department and University Honors

Advisor: Erik Learned-Miller

EMPLOYMENT **DeepMind, London** → **Google DeepMind, San Francisco**

Research Scientist → *Senior Research Scientist*

January 2021 - Present

Deep learning and computer vision research and development on robust visual recognition, self-supervised learning, attention architectures, and representation learning for remote sensing.

Adobe, Cambridge MA

Research Scientist

September 2019 - December 2020

Deep learning and computer vision research and development on adaptation to shifts (like weather), anytime inference, and image processing for recognition.

MIT, Cambridge MA

Visiting Research Scientist

January 2020 - December 2020

Guest lecturing, group discussions, and 1:1 student mentoring in the Freeman and Isola groups.

University of California, Berkeley

Graduate Student Researcher

September 2012 - May 2019

Thesis work on the end-to-end learning of image-to-image tasks, in particular the design and analysis of fully convolutional networks: multi-layer, multi-scale models for local recognition. Lead developer of Caffe, the open-source deep learning framework (caffe.berkeleyvision.org).

Facebook AI Research, Seattle

PhD Intern Researcher

July - October 2018

Visual recognition research supervised by Ross Girshick.

DeepMind, London

PhD Intern Researcher

June - September 2017

Deep reinforcement learning research supervised by Max Jaderberg.

OpenAI, San Francisco

PhD Intern Researcher

September 2016 - May 2017

Reinforcement learning research and engineering supervised by John Schulman. (20% time.)

AWARDS &
RECOGNITION

Mark Everingham Prize for service to the computer vision community through Caffe. (2017)
CVPR Best Paper Honorable Mention for fully convolutional networks. (2015)
ACM MM Open Source Award for the open source contribution of Caffe. (2014)
National Science Foundation Graduate Research Fellowship (2012-2015)
AI2000 Scholar Award for top 100 vision researchers as ranked by AMiner. (2020, 2022, 2023)
FADEx AI Fellow one of ten US representatives for US-France summit. (2018)

UC Berkeley EECS Chair Excellence Award department award for PhD students. (2012)
UMass Amherst CS Award department award for top undergraduate student. (2012)
ACM/UPE Award national award for achievement and outreach; one of four. (2011)
C.D. Youngren Research Award UMass Amherst research award & stipend; sole winner. (2010)
Commonwealth Honors UMass Amherst honors degree program. (2008-2012)

Reviewing Awards CVPR (2019, 2020, 2022), ICML (2020, 2021, 2022), ECCV (2020), NeurIPS (2020), ICLR (2022). TMLR Expert Reviewer (2023).

PUBLICATIONS

(* equal contribution, † equal advising)

Adaptive Randomized Smoothing: Certifying Multi-Step Defences against Adversarial Examples *NeurIPS 2024*.

S. Lyu*, S. Shaikh*, F. Shpilevskiy*, E. Shelhamer, M. Lecuyer. *Spotlight*.

Learned Embedding Fields for Multi-Source, Multi-Temporal Earth Observation Imagery *ICLRW 2024*.

C. Brown*, M. Kazmierski*, W. Rucklidge, V. Pasquarella, S. Alj, E. Schechter, S. Askay, A. Boukouvalas, E. Shelhamer.

Back to the Source: Diffusion-Driven Adaptation to Test-Time Corruptions *CVPR 2023*.

J. Gao*, J. Zhang*, X. Liu, T. Darrell, E. Shelhamer†, D. Wang†.

Seasoning Model Soups for Robustness to Adversarial and Natural Distribution Shifts *CVPR 2023*.

F. Croce, S.A. Rebuffi, E. Shelhamer, S. Goyal.

Adversarially Self-Supervised Pre-training Improves Accuracy and Robustness

ICLRW 2023.

S.A. Rebuffi, S. Goyal, O. Wiles, E. Shelhamer, S. Goyal.

Object Discovery and Representation Networks *ECCV 2022*.

O. Henaff, S. Koppula, E. Shelhamer, D. Zoran, A. Jaegle, A. Zisserman, J. Carreira, R. Arandjelović.

Evaluating the Adversarial Robustness of Test-Time Adaptive Defenses *ICML 2022*.

F. Croce*, S. Goyal*, T. Brunner*, E. Shelhamer*, T. Cemgil, M. Hein.

Anytime Dense Prediction with Confidence Adaptivity *ICLR 2022*.

Z. Liu, Z. Xu, H.J. Wang, T. Darrell, E. Shelhamer.

Perceiver IO: A General Architecture for Structured Inputs & Outputs *ICLR 2022*.

A. Jaegle, S. Borgeaud, J.B. Alayrac, C. Doersch, C. Ionescu, D. Ding, S. Koppula, D. Zoran, A. Brock, E. Shelhamer, O. Hénaff, M. Botvinick, A. Zisserman, O. Vinyals, J. Carreira.

Exploring Simple and Transferable Recognition-Aware Image Processing. *PAMI 2022*.

Z. Liu, H. Wang, T. Zhou, Z. Shen, B. Kang, E. Shelhamer, T. Darrell.

Where Should I Spend My FLOPS? Efficiency Evaluations of Visual Pre-training Methods *NeurIPSW 2022*.

S. Koppula, Y. Li, E. Shelhamer, A. Jaegle, N. Parthasarathy, R. Arandjelovic, J. Carreira, O. Henaff

Fighting Gradients with Gradients: Dynamic Defenses against Adaptive Attacks *ICLRW 2021*.

D. Wang, A. Ju, E. Shelhamer, D. Wagner, T. Darrell.

Tent: Fully Test-time Adaptation by Entropy Minimization *ICLR 2021*.

D. Wang*, E. Shelhamer*, S. Liu, B. Olshausen, T. Darrell. *Spotlight*.

Blurring the Line between Structure and Learning to Optimize and Adapt Receptive Fields *ICLRW 2019*.

E. Shelhamer, D. Wang, T. Darrell.

Infinite Mixture Prototypes *ICML 2019*

K. Allen, E. Shelhamer*, H. Shin*, J. Tenenbaum.

Conditional Networks for Few-Shot Semantic Segmentation *ICLRW 2018*

K. Rakelly*, E. Shelhamer*, T. Darrell, A. A. Efros, S. Levine.

Zero-Shot Visual Imitation *ICLR 2018*

D. Pathak*, P. Mahmoudieh*, M. Luo*, P. Agrawal*, D. Chen, F. Shentu, E. Shelhamer, J. Malik, A. A. Efros, T. Darrell. *Oral*.

Deep Layer Aggregation *CVPR 2018*

F. Yu, D. Wang, E. Shelhamer, T. Darrell. *Oral*.

Loss Is Its Own Reward: Self-Supervision for Reinforcement Learning *ICLRW 2017*

E. Shelhamer, P. Mahmoudieh, M. Argus, T. Darrell.

Clockwork Convnets for Video Semantic Segmentation *ECCVW 2016*

E. Shelhamer*, K. Rakelly*, J. Hoffman*, T. Darrell.

Fully Convolutional Multi-class Multiple Instance Learning *ICLRW 2015*

D. Pathak, E. Shelhamer, J. Long, T. Darrell.

Fully Convolutional Networks for Semantic Segmentation *PAMI 2016*

E. Shelhamer*, J. Long*, T. Darrell.

Fully Convolutional Networks for Semantic Segmentation *CVPR 2015*

J. Long*, E. Shelhamer*, T. Darrell. *Honorable Mention for Best Paper*. *Oral*.

Caffe: Convolutional Architecture for Fast Feature Embedding *ACM MM 2014*

Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, T. Darrell. *Open source award*.

cuDNN: Efficient Primitives for Deep Learning *NeurIPS DL Workshop 2014*

S. Chetlur, C. Woolley, P. Vandermersch, J. Cohen, J. Tran, B. Catanzaro, E. Shelhamer. *Oral*.

Communal cuts: sharing cuts across images *NeurIPS DISCOPT Workshop 2014*

E. Shelhamer, S. Jegelka, T. Darrell.

TALKS

Testing the Empirical and Certified Robustness of Test-Time Adaptation 2024

Invited talk at KAUST, Saudi Arabia.

Local and Adaptive Visual Recognition 2023

Invited Talk at AI2, Seattle.

Local and Adaptive Visual Recognition 2023

Invited Talk at UBC, Vancouver.

Local and Adaptive Visual Recognition 2023

Invited Talk at McGill, Montreal.

Local and Adaptive Visual Recognition 2023

Invited Talk at CMU Robotics Institute & NREC.

Local and Adaptive Visual Recognition 2023

Invited Talk at UMass Amherst, Amherst.

Test-Time Adaptation of Models and Inputs: Which and When? 2022

Spotlight talks at ICML workshops on Updatable ML and Dynamic Networks.

Robust Visual Recognition by Test-Time Adaptation 2022

Invited talk at CMU.

Robust Visual Recognition by Test-Time Adaptation 2022

Invited talk at Tidal (Google X).

Robust Visual Recognition by Test-Time Adaptation 2022

Keynote talk at DataFest Yerevan (industry conference in Armenia).

Fully Test-Time Adaptation and Defense by Entropy Minimization 2021

Invited talk at U. Tuebingen for Bethge and Brendel groups.

Test-time Optimization for Adaptive and Interactive Inference 2020

Invited talk at MIT for Isola group.

Blurring Structure and Learning for Adaptive Local Recognition 2019 - 2020

Invited talks at MIT, UBC, Adobe, DeepMind, BU, and NYU.

Few-Shot Segmentation by Guided Networks 2018

Invited talk for FADEX AI (exchange hosted by INRIA Sophia Antipolis, Grenoble, and Paris).

Pixels In, Pixels Out: Learning and Inference for Image-to-Image Tasks 2018

Invited talk at BarCamp Yerevan (industry conference in Armenia).

Clockwork Convolutional Networks for Video Segmentation 2017

Invited talk for the vision group at DeepMind, London.

End-to-End Recognition for Autonomous Vehicles 2016

Invited talk at Intelligent Vehicles in Gothenburg, Sweden.

Fully Convolutional Networks for Semantic Segmentation 2015–2016

Invited talks at MIT, Stanford, INRIA Paris, and UMass Amherst.

Caffe: Community Architecture for Fast Feature Embedding *July 2015*
Workshop talk at ICML MLOSS.

DIY Deep Learning for Vision: a Hands-On Tutorial *2014–2016*
CVPR 2015 conference tutorial. Attended by 400+.
ECCV 2014 conference tutorial. Attended by 100+.
GTC 2015–2016 industry conference tutorials. Attended by 200+.

The Caffe Latest Roast *2014–2016*
Invited industry talks at Sony, Yahoo! Japan, Samsung, Amazon & A9.

Demos

Interactive Mapping of the Earth with Embedding Fields *CVPR Demos 2024*.
C. Brown*, Sean Askay*, M. Kazmierski, W. Rucklidge, V. Pasquarella, S. Alj, E. Schechter, A. Boukouvalas, E. Shelhamer.

Back to the Source: Adapting Inputs by Diffusion to Counter Corruptions *CVPR Demos 2023*.
J. Gao*, J. Zhang*, X. Liu, T. Darrell, E. Shelhamer[†], D. Wang[†].

Caffe: Image Classification with DIY Deep Learning *ICVSS 2014 Poster & Demo*.
Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, T. Darrell.

SERVICE

Area Chair: CVPR (2021, 2023, 2024), ICCV (2021, 2023), ECCV (2024), NeurIPS (2023, 2024), ICLR (2024), ICML (2024)

Editor: TMLR Action Editor (2023-)

Reviewer: CVPR (2014-2020), ICCV (2015-2019), ECCV (2016-2021), NeurIPS (2014-2022), ICML (2015-2023), ICLR (2015-2023), JMLR (2020-), PAMI (2016-) and TMLR (2022-2023)

Workshops:

- lead organizer of the 1st Workshop on Test-Time Adaptation at CVPR 2024

Tutorials:

- lead organizer of ML for Remote Sensing (in submission for NeurIPS 2024)
- lead organizer of DIY Deep Learning at CVPR 2015
- lead organizer of DIY Deep Learning at ECCV 2014

Graduate Admissions: reviewer for PhD applications in AI at UC Berkeley. *2016-2017*

CSGSA: officer of the UC Berkeley Computer Science Graduate Association. *2013-2015*

UMass Amherst ACM: president. *2011-2012*

UMass Amherst Deans' Committee: volunteer for science outreach & open houses. *2010-2012*

ADVISING & MENTORING

Ph.D Committees

Dequan Wang, Computer Science, UC Berkeley 2022

Motasem Alfara, Electrical and Computer Engineering, KAUST 2024

Wonho Bae, Computer Science, UBC 2025 (expected)

Ph.D Student Mentoring

Spandan Madan (2024-), PhD student at Harvard.

Fahim Tajwar (2023-2024), PhD student at CMU.

Evgenia Rusak (2023-2024), PhD student at U. Tübingen.

Dequan Wang (2018-2022), PhD student, then professor at SJTU.

Zhuang Liu (2020-2022), PhD student, then research scientist at Meta.

Parsa Mahmoudieh (2016-2021), PhD student, then research engineer at DeepMind.

Kate Rakelly (2016-2019), PhD student, then research scientist at Cruise.

M.S. Student Mentoring

Isaac Akintaro (2024), University of Birmingham, then PhD student at U. of Birmingham.
Ilakya Prabhakar (2022-2023), University of Edinburgh, then researcher at Alan Institute.
Max Argus (2016-2019), University of Freiburg, then postdoc at University of Freiburg.

B.S. Students Mentoring

Cullen Anderson (2023-2024), undergraduate student at UMass Amherst.

Industry Mentoring

Nikita Saxena (2023-), student researcher at Google Bangalore, then research engineer at Google DeepMind. Lukas Hoyer (2024-), software engineer at Google Zurich.

DeepMind Scientific Advising (2023-) for 1:1 support of research and engineering roles.

DeepMind Scholar Mentorship (2022-) for 1:1 support of MS and PhD students.

Mentoring Events: in-person sessions at conferences (CVPR'24, ICLR'24, ECCV'24, CVPR'23), virtual sessions at conferences (ICML'20, ICLR'20), virtual session at summer school (EEML'22).

TEACHING

Graduate Student Instructor (TA)

University of California, Berkeley

Fall 2014

DIY Deep Learning: made tutorials, led seminars, and answered 1000+ questions online.

Graduate Student Instructor (TA)

University of California, Berkeley

Fall 2013

CS 188 Artificial Intelligence: taught by Dan Klein and Pieter Abbeel and offered to four hundred Berkeley undergraduates as well as online to thousands. Led discussion section 3 hours/week, held office hours, created exams and tutorials, and helped maintain the project grading system.

Guest Lecturer

Model Soups for Efficient Ensembling and Adaptation at KAUST for Prof. El-Hosseiny, 2024

Test-Time Adaptation for Visual Recognition at KAUST for Prof. El-Hosseiny, 2021

Weaving Deep Networks at MIT for Prof. Isola, 2019

Adaptive Local Recognition at UMass Amherst for Prof. Learned-Miller, 2019

Image-to-Image Learning and Inference at Yerevan State U. for Instructor Khachatryan, 2018

Image-to-Image Learning and Inference at UC Berkeley for Profs. Darrell and Prof. Efros, 2015

DIY Deep Learning with Caffe at UC Berkeley for Prof. Canny, 2015

DIY Deep Learning with Caffe at Stanford for Instructor Karpathy, 2015

DIY Deep Learning with Caffe at UMass Lowell for Prof. Saenko, 2014

Industry Tutorial Instructor

Embedded Vision Alliance

Summer & Fall 2016

Deep Learning for Vision with Convnets and Caffe: developed the tutorial content with slides, examples, and hands-on exercises then led the delivery of the tutorial in-person at events for software developers, hardware engineers, and image processing experts. The tutorial was held as a half-day event at the Embedded Vision Alliance summit in Santa Clara, CA and as a full-day event in Cambridge, MA for 100+ attendees each.