

# Laurens van der Maaten, Distinguished Research Scientist

---

lvdmaaten@gmail.com, <http://lvdmaaten.github.io>

## EXPERIENCE *Meta AI*

- 2023–now: Distinguished Research Scientist, GenAI Llama team.
  - With co-lead, responsible for technical direction of Meta’s Llama foundational modeling team of 100+ Research Scientists and Engineers.
  - Developed original plan for Llama 3, drove integration of vision and speech capabilities, made key contributions to pre-training and post-training execution, re-architected human evaluations, wrote paper, *etc.*
  - Developed agent research strategy for Llama team.
  - Leadership role in defining the Meta’s broader AI strategy, in close collaboration with company leadership.
- 2023: FAIR Accel lead, Senior Research Director.
  - Managed a team of 130+ Research Scientists, Research Engineers, Software Engineers, Designers, and UX Researchers across all levels of seniority (IC4 to IC9) and across six different sites in the US and Europe.
  - Supported all large, focused research projects (10+ researchers per project, multi-year efforts) in the FAIR team.
  - Responsible for all key processes including hiring, performance evaluation, people management, and budget.
- 2021–2022: FAIR AI for Math & Science lead, Research Director.
  - Managed a team of 60+ Research Scientists and Research Engineers, across all levels of seniority (IC4 to IC9) and across six different sites in the US and Europe.
  - Pioneered culture and processes that enable large, focused research projects (10+ researchers per project, multi-year efforts) in a bottom-up research organization. Strengthened FAIR’s communication around major research results.
  - Responsible for all key processes including hiring, performance evaluation, people management, and budget.
  - Team successes include development of the ESM protein models and ESM-Atlas, the Cicero full-press Diplomacy bot, the Galileo scientific language model, and the current state-of-the-art in automatic theorem proving.
  - Initiated several new research efforts at the intersection of AI and climate technology and the intersection of AI and augmented-reality technology.
- 2019–2021: FAIR New York site lead, Research Director (since 2020).
  - Managed all 30+ FAIR Research Scientists based in New York, across all levels of seniority (IC4 to IC9).
  - Closely partnered with Research Engineering organization of 40+ engineers.
  - Responsible for all key processes including hiring, performance evaluation, and people management.

- Developed several new processes and contributed to key strategic initiatives. Examples include development of FAIR’s publication review process and of clear career paths for senior researchers.
- Contributed to several other key research innovations, including the development of Fisher information loss as a versatile alternative for differential privacy (UAI 2021 Best Paper Award).
- 2015–2019: Research Scientist, Fundamental AI Research (FAIR).
  - Pioneered web-scale weakly supervised training of image-recognition models at Meta. These models were presented by Meta’s CTO at the F8 developer conference in 2018, and form the backbone of all major image-recognition systems used in Facebook and Instagram.
  - Developed the CrypTen framework that allows developers to execute PyTorch code on data that is secret-shared across parties, enabling fully private training and inference of deep-learning models.
  - Contributed to several other key research innovations, including the DenseNet convolutional network architecture (CVPR 2017 Best Paper Award) and the first models that learned to associate images and language from web data.

*Delft University of Technology.* Assistant Professor (with tenure). Delft, The Netherlands, January 2015–February 2018.

*Delft University of Technology.* Assistant Professor (tenure-track). Delft, The Netherlands, April 2013–December 2014.

*Netherlands Institute for Advanced Study.* Fellow. Wassenaar, The Netherlands, February 2013–April 2013.

*Delft University of Technology.* Post-doc in (EU FP7) SSPNet project, working with Dr. Emile Hendriks, Dr. Bob Duin, Dr. Marco Loog, and Dr. David Tax. Delft, The Netherlands, August 2009–March 2013.

*University of California, San Diego.* Post-doc with Prof. Lawrence Saul. San Diego CA, December 2009–November 2010.

*Tilburg University.* Ph.D. student under supervision of Prof. Eric Postma and Prof. Jaap van den Herik, Tilburg, The Netherlands, October 2008–August 2009. Ph.D. received at June 23<sup>rd</sup> 2009 (*cum laude*).

*Dutch State Service for Cultural Heritage (RCE).* Ph.D. student working with Dr. Guus Lange, Amersfoort, The Netherlands, August 2005–August 2009.

*Maastricht University.* Ph.D. student under supervision of Prof. Eric Postma and Prof. Jaap van den Herik, Maastricht, The Netherlands, August 2005–October 2008.

## RESEARCH VISITS

*University of Geneva.* Visiting postdoctoral researcher working with Dr. Marc Mehu, Geneva, Switzerland, January 2012.

*Imperial College London.* Visiting postdoctoral researcher working with Prof. Maja Pantic, London, UK, July 2011.

*University of Toronto.* Visiting Ph.D. student working with Prof. Geoffrey Hinton, Toronto, ON, Canada, January 2008–June 2008.

**EDUCATION**

*Tilburg University.* Doctoral degree (*cum laude*) for Ph.D. thesis “Feature Extraction from Visual Data”. Thesis defense on June 23<sup>rd</sup> 2009.

*Maastricht University.* Doctorandus (= Dutch M.Sc.) in Knowledge Engineering and Computer Science, Maastricht, The Netherlands, September 2001–August 2005.

**RESEARCH INTERESTS**

I am broadly interested in machine learning and computer vision research. I have worked on topics including dimensionality reduction, embedding, metric learning, generative models, deep learning, privacy-preserving machine learning, large-scale weakly supervised learning, self-supervised learning, and applications of machine learning in the arts and sciences.

**SELECTED PUBLICATIONS**

My *h*-index is 53 and my total number of citations is 105,045 (Google Scholar; retrieved November, 2023). The following publications are representative for my work:

Llama Team\*, AI @ Meta. The Llama 3 Herd of Models. In *arXiv 2407.21783*, 2024. \*I was one of the two main tech leads for Llama 3.

A. Hannun, C. Guo, and L.J.P. van der Maaten. Measuring Data Leakage in Machine-Learning Models with Fisher Information. In *Proceedings of Uncertainty in Artificial Intelligence (UAI)*, 2021. [Acceptance rate: 32.3%; 31 citations; Best Paper Award]

B. Knott, S. Venkataraman, A.Y. Hannun, S. Sengupta, M. Ibrahim, and L.J.P. van der Maaten. CryptTen: Secure Multi-Party Computation Meets Machine Learning. In *Advances of Neural Information Processing Systems (NeurIPS)*, 2021. [Acceptance rate: 26%; 212 citations].

D. Mahajan, R.B. Girshick, V. Ramanathan, K. He, M. Paluri, Y. Li, A. Bharambe, and L.J.P. van der Maaten. Exploring the Limits of Weakly Supervised Pretraining. In *European Conference on Computer Vision (ECCV)*, pages 185–201, 2018. [Acceptance rate: 31.8%; 1,386 citations].

G. Huang, Z. Liu, K.Q. Weinberger, and L.J.P. van der Maaten. Densely Connected Convolutional Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [Best paper award; Oral acceptance rate: 2.65%; 40,481 citations].

L.J.P. van der Maaten. Accelerating t-SNE using Tree-Based Algorithms. *Journal of Machine Learning Research* 15(Oct):3221–3245, 2014. [Impact factor: 3.42; 2,740 citations]

L.J.P. van der Maaten and G.E. Hinton. Visualizing High-Dimensional Data Using t-SNE. *Journal of Machine Learning Research* 9(Nov):2579–2605, 2008. [Impact factor: 2.949; 38,711 citations]

**AWARDS**

*UAI Best Paper Award.* Best paper award at the Uncertainty in Artificial Intelligence (UAI) conference for paper: A. Hannun, C. Guo, and L.J.P. van der Maaten. Measuring Data Leakage in Machine-Learning Models with Fisher Information, 2021.

*CVPR Best Paper Award.* Best paper award at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) for paper: G. Huang, Z. Liu, K.Q. Weinberger, and L.J.P. van der Maaten. Densely Connected Convolutional Networks, 2017.

*SNN Machine Learning Award.* Awarded once per year by the Netherlands Society for Neural Networks (SNN) to a researcher or research group who has achieved an important result in machine learning. I received the award for my work on new machine-learning techniques for data visualization. Amount: 5k€.

*Merck Molecular Activity Visualization Challenge Winner.* Best visualization of molecular activity data of Merck published on Kaggle.com. Amount: 2k\$.

*NIAS-Lorentz Fellowship.* Awarded by the Netherlands Institute for Advanced Study and the Lorentz Center for my research on “Texture Analysis for the Arts”. Awarded in March 2012. Amount: 9k€.

*AAAI-08 Most Innovative Video Award.* Awarded for the video entitled “Digital Analysis of Van Gogh Paintings” at the conference of the Association for the Advancement of Artificial Intelligence, July 2008, Chicago, IL, USA. Joint work with Eric Postma.

*MUSCLE CIS Coin Competition 2007 Winner.* Best classification performance in a benchmark on automatic classification of modern coins based on coin photographs, Vienna, Austria, August 2007. Joint work with Paul Boon.

*Travel awards.* I received travel awards for NIPS 2010 (sponsored by Google; \$800), AI-STATS 2009 (sponsored by Nokia; \$550), and CAA 2009 (sponsored by Computer Applications for Archaeology; \$400).

## **SOFTWARE**

I am the lead developer of CrypTen. I was one of the developers of Classy Vision and of the visdom tool for scientific visualization (over 6,900 stars on Github). I maintain the Matlab Toolbox for Dimensionality Reduction. Since its introduction in 2006, the toolbox has been downloaded over 40,000 times. I also maintain a website on t-Distributed Stochastic Neighbor Embedding (t-SNE), which is a new technique for the visualization of high-dimensional data that I developed together with Prof. Geoffrey Hinton. My website on t-SNE contains implementations of the technique in various programming language. My personal homepage also features Matlab implementations of, among others, state-of-the-art system for writer identification, fields of experts, and matrix relational embedding. I also contributed to the open-source visualization software Divvy and to several (Py)Torch repositories.

## **RESEARCH GRANTS**

### **Learning the Fundamental Symmetries in Video Data.**

*Funding agency:* Netherlands Scientific Organisation (NWO)

*PI:* Prof.dr. M. Welling (University of Amsterdam)

*Co-PI:* Dr. L.J.P. van der Maaten (TU Delft)

*Amount:* 217k€

*Start date / end date:* January 2015 / January 2019

*Duration:* 48 months

### **HDCLUSTER**

Development of novel clustering methods for gene expression, proteomic and imaging readouts by concurrently mining the Allen Brain Atlas.

*Funding agency:* FET Flagship (Human Brain Project)

*PIs:* Prof.dr. B. Lelieveldt (LUMC), Prof.dr. M.J.T. Reinders (TU Delft

/ LUMC), Dr. Laurens van der Maaten (TU Delft / LUMC), Prof. A. van den

Maagdenberg (LUMC), Dr. W. van Roon (LUMC)

*Amount:* 234k€

*Start date / end date:* September 2014 / September 2016.

*Duration:* 24 months

**Technology in Motion.**

*Funding agency:* Netherlands Scientific Organisation (NWO)

*Project coordinator:* Prof.dr. B. van Hilten (LUMC)

*Partners:* Leiden University Medical Center, TU Delft (EWI and TBM), Free University Amsterdam.

*PI at TU Delft-EWI:* Dr. L.J.P. van der Maaten

*Amount for TU Delft-EWI:* 210k€

*Start date / end date:* January 2014 / January 2018

*Duration:* 48 months

**Learning from Corruption.**

*Funding agency:* Netherlands Scientific Organisation (NWO)

*PI:* Dr. L.J.P. van der Maaten

*Amount:* 207k€

*Start date / end date:* July 2014 / July 2018

*Duration:* 48 months

**Integration of Technological Solutions for Imaging, Detection, and Digitisation of Hidden Elements in Artworks (INSIDDE).**

*Funding agency:* European Commission (FP7-ICT 2011.8.2; STReP)

*Project coordinator:* Dr. S. Garcia (Treelogic)

*PI at TU Delft:* Dr. L.J.P. van der Maaten

*Partners:* Treelogic (ES), Universidad de Oviedo (ES), Fundacion ITMA (ES), TU Delft (NL), 3DDynamics (BE), Consiglio Nazionale delle Ricerche (IT), Regional Museum of History - Stara Zagora (BG), and Doerner Institute (DE).

*Amount for TU Delft:* 401k€

*Start date / end date:* January 2013 / January 2016

*Duration:* 36 months

**Smart Assisted Living involving Informal careGivers (SALIG++).**

*Funding agency:* European Commission and Netherlands Organisation for Health Research and Development (ZonMW); Ambient Assisted Living Programme

*Project coordinator:* Dr. G. Juell-Skielse (Stockholm University)

*PI at TU Delft:* Dr. L.J.P. van der Maaten

*Partners:* Stockholm University (SE), Almende (NL), HI Iberia (ES), TU Delft (NL), FeSalud (ES), Stockholm LänsLandsting (SE), PIAP (PL), Actimage (LU), and Divitel (NL).

*Amount for TU Delft:* 280k€

*Start date / end date:* June 2013 / June 2016

*Duration:* 36 months

**Visualizing High-Dimensional Data using Dimensionality Reduction.**

*Funding agency:* Netherlands Scientific Organisation (NWO)

*PI:* Dr. L.J.P. van der Maaten

*Amount:* 28k€

*Start date / end date:* December 2010 / December 2011

*Duration:* 12 months

*Research team at UC San Diego:* Dr. L.J.P. van der Maaten and Prof. L.K. Saul

**TEACHING &  
SUPERVISION**

I have received the Dutch University Qualification Teaching (UTQ) in October 2012.

At Meta, I have (co-)supervised the following interns and residents:

- Yun-Yun Tsai, 2024.
- Basile van Hoorick, 2023.
- Steffen Schneider, 2022.
- Amrita Roy Chowdhury, 2021.
- Ruihan Wu, 2021.
- Karan Desai, 2021.
- Tony Ginart, 2021.
- Mimee Xu, 2020.
- Kexin Yi, 2020.
- Eltayeb Ahmed, 2020.
- Yin Cui, 2019.
- Chuan Guo, 2019 and 2017.
- Terrence DeVries, 2019.
- Mandela Patrick, 2018.
- Abhimanyu Dubey, 2018.
- Hexiang Hu, 2018.
- Cihang Xie, 2018.
- Ishan Misra, 2017.
- Justin Johnson, 2017.
- Andreas Veit, 2017.
- Haytham Fayek, 2016.
- Ang Li, 2016.
- Ranjay Krishna, 2016.

I have supervised two post-doctoral researchers:

- Hamdi Dibeklioglu. *Machine Learning and Computer Vision for Painting Analysis*. July 2013–January 2015.
- Julian Kooij. *Automatic Analysis of Tremors of Parkinson’s Patients*. Started October 2014–January 2015.

The following Ph.D. students have completed their thesis under my supervision:

- Gorkem Saygili. *Depth Estimation by Combining Active and Passive Stereo Techniques*. Defense date: June 17, 2015. Co-supervised with Dr. Emile Hendriks.
- Wouter Kouw. *Machine Learning and Computer Vision for Painting Analysis*. Defense date: June 5, 2018. Co-supervised with Dr. Marco Loog.
- Wenjie Pei. *Machine Learning and Computer Vision for Assistance of Elderly People*. Defense date: June 13, 2018. Co-supervised with Dr. David Tax.
- Taygun Kecec. *Learning by Marginalizing Corrupted Features*. Defense date: June 7, 2019. Co-supervised with Dr. David Tax.
- Lu Zhang. *Tracking for Automatic Analysis of Human Eating Behavior*. Defense date: January 30, 2015.

I was member of the thesis committee of the following Ph.D. students:

- Damith Senanayake, *Exploratory Analysis of Highly Dimensional Data: Parametric Methods for Dimensionality Reduction, Visualization and Feature Extraction with Applications in Computational Biology*, University of Melbourne, April 2020.
- Stephen Tyree, *Some Approaches to Parallel and Distributed Machine Learning*, Washington University in St. Louis, November 2014.
- Philemon Brakel, *The Estimation and Application of Unnormalized Statistical Models*, Ghent University, July 2014.

## TEACHING

I have been a (guest) lecturer of the following M.Sc./B.Sc. courses:

- Introduction to Computer Vision (African Master on Machine Intelligence; M.Sc.), 2019, 2020, 2021, 2022, 2023, 2024
- Convolutional Networks (Georgia Tech; B.Sc.), 2022
- Big Ideas in AI (New York University; B.Sc.), 2021
- Machine Learning (Cornell Tech; M.Sc.), 2015, 2016, 2017
- Computer Vision (Delft University of Technology; M.Sc.) 2013, 2014
- Machine Learning (Delft University of Technology; M.Sc.) 2013, 2014
- Datamining (Delft University of Technology; B.Sc.) 2014
- Imaging Project (Delft University of Technology; B.Sc.) 2012
- Bachelor Seminar (Delft University of Technology; B.Sc.) 2011, 2012

I gave lectures at the following Ph.D.-level summer schools:

- Advanced Pattern Recognition (ASCI), Delft, The Netherlands, 2011, 2012, 2013
- Computer Vision by Learning (ASCI), Amsterdam, The Netherlands, 2014, 2016, 2018
- Computational Intelligence (SIKS), Utrecht, The Netherlands, 2011
- Summer School on Social Signal Processing, Mullsjö, Sweden, 2013
- Spring School on Social Interaction Computing, Almen, The Netherlands, 2010
- Datamining Summer School, Maastricht, The Netherlands, 2008

## ACTIVITIES

Together with Aaron Adcock, Deepti Ghadiyaram, Angelina Wang, Vicente Ordonez-Roman, Judy Hoffman, Reza Shokri, Cristian Canton Ferrer, and Been Kim, I co-organized the ECCV 2022 Workshop on Responsible Computer Vision.

Together with Deepti Ghadiyaram, Abhimanyu Dubey, Han Zhao, Dhruv Mahajan, and Olga Russakovsky, I co-organized the CVPR 2021 Workshop on Responsible Computer Vision.

Together with Alex Berg, Cheng-Fang Yu, Dhruv Mahajan, Ishan Misra, Vignesh Ramanathan, and Zhenheng Yang, I co-organized the CVPR 2021 Workshop on Extreme Computer Vision.

Together with Pengtao Xie, Shanghang Han, Pulkit Agrawal, Ishan Misra, Cynthia Rudin, Abdelrahman Mohamed, Wenzhen Yuan, Barret Zoph, Xingyi Yang, and Eric Xing, I co-organized the NeurIPS 2020 Workshop on Self-Supervised Learning.

Together with Vignesh Ramanathan, Dhruv Mahajan, Alex Berg, Ishan Misra, and Rahul Sukthankar, I co-organized the ICCV 2019 Workshop on Extreme Vision Modeling.

Together with Dr. Marco Loog (Delft University of Technology), I organized the Belgian-Dutch Machine Learning Conference (Benelearn) 2015.

In 2014, I served as a young expert for the “National Think Tank” (*Nationale Denktank*) on Big Data.

Together with Dr. Cees Snoek (University of Amsterdam), I have organized the inaugural Netherlands Conference on Computer Vision (NCCV) 2014.

With Prof. Robert Erdmann, I have organized a NIAS Workshop on “*New Technologies for Art Conservation and History*” (2013).

I have co-organized the Advances in Neural Information Processing Systems (NIPS) 2010 Workshop on “*Challenges of Data Visualization*” (with Prof. Barbara Hammer, Prof. Fei Sha, and Prof. Alexander Smola).

## SERVICE

- Neural Information Processing Systems (NeurIPS): Area Chair in 2015 and 2018; Senior Area Chair in 2020, 2021, 2022, and 2023; Best Paper Award Committee in 2020.
- International Conference on Machine Learning (ICML): Area Chair in 2016, 2017, 2018, 2019, 2021, and 2024; Best Paper Award Committee in 2022; Tutorial Chair in 2022.
- IEEE Conference on Computer Vision and Pattern Recognition (CVPR): Area Chair in 2018 and 2021; Senior Area Chair in 2022 and 2023; Best Paper Award Committee in 2023.
- International Conference on Learning Representations (ICLR): Area Chair in 2017, 2018, and 2022.
- European Conference on Computer Vision (ECCV): Area Chair in 2016 and 2024.

From 2016 to 2020, I was an associate editor for IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI). In 2014, I was associate editor of IEEE Transactions on Neural Networks and Learning Systems (TNNLS).

I have reviewed research proposals of the Swiss National Science Foundation (SNSF), the Belgian Agency for Innovation by Science and Technology (IWT), and the European Commission.

From 2012 to 2014, I was a book reviewer for Zentralblatt MATH.

I have reviewed papers for over 35 different journals. Five prominent journals for which I served as a reviewer are:

- Journal of Machine Learning Research
- Machine Learning
- Proceedings of the National Academy of Science
- Science
- Neural Computation



I have been a reviewer for over 27 different conferences. Five prominent conferences for which I served as a reviewer are:

- Advances in Neural Information Processing Systems (NIPS) 2010, 2012, 2013, 2014, 2015, 2016, 2017
- International Conference on Machine Learning (ICML) 2012, 2013, 2014, 2015, 2016, 2017
- Artificial Intelligence and Statistics (AI-STATS) 2010, 2011, 2012, 2013, 2014, 2015, 2016
- ACM SIGKDD Conf. on Knowledge Discovery and Data Mining (KDD) 2011, 2015
- IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2015, 2016, 2017

**CONFERENCE PUBLICATIONS** V.V. Ramaswamy, S.Y. Lin, D. Zhao, A.B. Adcock, L.J.P. van der Maaten, D. Ghadiyaram, and O. Russakovsky. Beyond web-scraping: Crowd-sourcing a geographically diverse image dataset. To appear in *Advances of Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks*, 2023.

K. Desai, I. Misra, J.C. Johnson, and L.J.P. van der Maaten. Scaling up Instance Segmentation using Approximately Localized Phrases. In *British Machine Vision Conference (BMVC)*, 2022.

A. Roy Chowdhury, C. Guo, S. Jha, and L.J.P. van der Maaten. EIFFeL: Ensuring Integrity for Federated Learning. To appear in *ACM Conference on Computer and Communications Security (ACM CCS)*, 2022. [Acceptance rate: 18.0%; 21 citations]

C. Guo, B. Karrer, K. Chaudhuri, and L.J.P. van der Maaten. Bounding Training Data Reconstruction in Private (Deep) Learning. In *International Conference on Machine Learning (ICML)*, 2022. [Acceptance rate: 21.9%; 27 citations]

M. Hall, L.J.P. van der Maaten, L. Gustafson, and A. Adcock. A Systematic Study of Bias Amplification. To appear in *NeurIPS 2022 Workshop on Trustworthy and Socially Responsible Machine Learning*, 2022. [33 citations]

M. Singh, L. Gustafson, A. Adcock, V. de Freitas Reis, B. Gedik, R. Prateek Kosaraju, D. Mahajan, R.B. Girshick, P. Dollar, and L.J.P. van der Maaten. Revisiting Weakly Supervised Pre-Training of Visual Perception Models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [Acceptance rate: 25.3%; 25 citations]

R. Girdhar\*, M. Singh\*, N. Ravi\*, L.J.P. van der Maaten, A. Joulin, and I. Misra\*. Omnivore: A Single Model for Many Visual Modalities. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [Acceptance rate: 25.3%; 116 citations]

M. Xu, L.J.P. van der Maaten, A.Y. Hannun. Data Appraisal Without Data Sharing. In *Proceedings of Artificial Intelligence and Statistics Conference (AI-STATS)*, 2022. [Acceptance rate: 29.2%]

B. Knott, S. Venkataraman, A.Y. Hannun, S. Sengupta, M. Ibrahim, and L.J.P. van der Maaten. CrypTen: Secure Multi-Party Computation Meets Machine Learning. In

*Advances of Neural Information Processing Systems (NeurIPS)*, 2021. [Acceptance rate: 25.7%; 212 citations]

R. Wu, C. Guo, A. Hannun, and L.J.P. van der Maaten. Fixes That Fail: Self-Defeating Improvements in Machine-Learning Systems. In *Advances of Neural Information Processing Systems (NeurIPS)*, 2021. [Acceptance rate: 25.7%]

A. Hannun, C. Guo, and L.J.P. van der Maaten. Measuring Data Leakage in Machine-Learning Models with Fisher Information. In *Proceedings of Uncertainty in Artificial Intelligence (UAI)*, 2021. [Acceptance rate: 32.3%; 31 citations; Best Paper Award]

R. Wu, C. Guo, F. Wu, R. Kidambi, L.J.P. van der Maaten, and K.Q. Weinberger. Making Paper Reviewing Robust to Bid Manipulation Attacks. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2021. [Acceptance rate: 21.75%; 21 citations]

I. Misra and L.J.P. van der Maaten. Self-supervised learning of pretext-invariant representations. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [Acceptance rate: 25%; 1,322 citations]

C. Guo, T. Goldstein, A.Y. Hannun, and L.J.P. van der Maaten. Certified data removal from machine learning models. In *International Conference on Machine Learning (ICML)*, 2020. [Acceptance rate: 21.8%; 218 citations]

A. Bakhtin, L.J.P. van der Maaten, J.C. Johnson, L. Gustafson, and R.B. Girshick. PHYRE: A New Benchmark for Physical Reasoning. In *Advances of Neural Information Processing (NeurIPS)*, 2019. [Acceptance rate: 21%; 103 citations]

T. DeVries\*, I. Misra\*, C. Wang\*, and L.J.P. van der Maaten. Does Object Recognition Work for Everyone? In *CVPR Workshop on Computer Vision for Global Challenges*, 2019. [219 citations]

H. Hu, I. Misra, and L.J.P. van der Maaten. Evaluating Text-to-Image Matching using Binary Image Selection (BISON). In *ICCV Workshop Closing the Loop Between Vision and Language*, 2019.

Y. Wang, Z. Lai, G. Huang, B.H. Wang, L.J.P. van der Maaten, M. Campbell, and K.Q. Weinberger. Anytime Stereo Image Depth Estimation on Mobile Devices. In *International Conference on Robotics and Automation (ICRA)*, 2019. [183 citations]

A. Dubey, L.J.P. van der Maaten, Z. Yalniz, Y. Li, and D. Mahajan. Defense Against Adversarial Images using Web-Scale Nearest-Neighbor Search. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [Acceptance rate: 25.2%; 62 citations]

C. Xie, Y. Wu, L.J.P. van der Maaten, A. Yuille, and K. He. Feature Denoising for Improving Adversarial Robustness. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [Acceptance rate: 25.2%; 884 citations]

D. Mahajan, R.B. Girshick, V. Ramanathan, K. He, M. Paluri, Y. Li, A. Bharambe, and L.J.P. van der Maaten. Exploring the Limits of Weakly Supervised Pretraining. In *European Conference on Computer Vision (ECCV)*, pages 185–201, 2018. [Acceptance rate: 31.8%; 1,386 citations].

- I. Misra, R.B. Girshick, R. Fergus, M. Hebert, A. Gupta, and L.J.P. van der Maaten. Learning by Asking Questions. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [Oral acceptance rate: 2.1%; 80 citations]
- A. Veit, M. Nickel, S. Belongie, and L.J.P. van der Maaten. Separating Self-Expression and Visual Content in Hashtag Supervision. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [Acceptance rate: 29.6%; 35 citations]
- G. Huang, S. Liu, L.J.P. van der Maaten, and K.Q. Weinberger. CondenseNet: An Efficient DenseNet using Learned Group Convolutions. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [Acceptance rate: 29.6%; 881 citations]
- B. Graham, M. Engelcke, and L.J.P. van der Maaten. 3D Semantic Segmentation with Submanifold Sparse Convolutional Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [Acceptance rate: 29.6%; 1,300 citations]
- C. Guo, M. Rana, M. Cisse, and L.J.P. van der Maaten. Countering Adversarial Images using Input Transformations. In *International Conference on Learning Representations (ICLR)*, 2018. [Acceptance rate: 33.6%; 1,361 citations]
- G. Huang, D. Chen, T. Li, F. Wu, L.J.P. van der Maaten, and K.Q. Weinberger. Multi-Scale Dense Convolutional Networks for Efficient Prediction. In *International Conference on Learning Representations (ICLR)*, 2018. [Oral acceptance rate: 2.5%; 833 citations]
- A. Li, A. Jabri, A. Joulin, and L.J.P. van der Maaten. Learning Visual N-Grams from Web Data. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2017. [Acceptance rate: 29.0%; 119 citations]
- J. Johnson, B. Hariharan, L.J.P. van der Maaten, J. Hoffman, L. Fei-Fei, C.L. Zitnick, and R.B. Girshick. Inferring and Executing Programs for Visual Reasoning. To appear in *Proceedings of the International Conference on Computer Vision (ICCV)*, 2017. [Oral acceptance rate: 2.1%; 569 citations]
- H. Fayek, L.J.P. van der Maaten, G. Romigh, and R. Mehra. On Data-Driven Approaches to Head-Related-Transfer Function Personalization. In *Audioengineering*, 2017.
- G. Huang, Z. Liu, K.Q. Weinberger, and L.J.P. van der Maaten. Densely Connected Convolutional Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [Best paper award; Oral acceptance rate: 2.65%; 40,481 citations]
- J. Johnson, B. Hariharan, L.J.P. van der Maaten, L. Fei-Fei, C.L. Zitnick, and R.B. Girshick. CLEVR: A Diagnostic Dataset for Compositional Language and Elementary Visual Reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [Acceptance rate: 29.0%; 2,070 citations]
- A. Joulin\*, L.J.P. van der Maaten\*, A. Jabri, and N. Vasilache. Learning Visual Features from Large Weakly Supervised Data. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 67–84, 2016. [Acceptance rate: 26.6%; 388 citations]
- A. Jabri, A. Joulin, and L.J.P. van der Maaten. Revisiting Visual Question Answering Baselines. In *Proceedings of the European Conference on Computer Vision (ECCV)*,

pages 727–739, 2016. [Acceptance rate: 26.6%; 305 citations]

R. Collobert, L.J.P. van der Maaten, and A. Joulin. Torchnet: An Open-Source Platform for (Deep) Learning Research. In *ICML Machine Learning Systems Workshop*, 2016.

L. Zhang, H. Dibeklioglu, and L.J.P. van der Maaten. Speeding Up Tracking by Ignoring Features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1266–1273, 2014. [Acceptance rate: 29.9%]

Y. Guo, H. Dibeklioglu, and L.J.P. van der Maaten. Graph-Based Kinship Recognition. In *Proceedings of the International Conference on Pattern Recognition (ICPR)*, pages 4287–4292, 2014. [86 citations]

L. Zhang and L.J.P. van der Maaten. Improving Object Tracking by Adapting Detectors. In *Proceedings of the International Conference on Pattern Recognition (ICPR)*, pages 1218–1223, 2014.

G. Saygili, L.J.P. van der Maaten, and E.A. Hendriks. Stereo Similarity Metric Fusion Using Stereo Confidence. In *Proceedings of the International Conference on Pattern Recognition (ICPR)*, pages 2161–2166, 2014.

G. Saygili, L.J.P. van der Maaten, and E.A. Hendriks. Hybrid Kinect Depth Map Refinement For Transparent Objects. In *Proceedings of the International Conference on Pattern Recognition (ICPR)*, pages 2751–2756, 2014.

P. Messier, C.R. Johnson, H. Wilhelm, W.A. Sethares, A.G. Klein, P. Abry, S. Jaffard, H. Wendt, S. Roux, N. Pustelni, N. van Noord, L.J.P. van der Maaten, and E.O. Postma. Automated Surface Texture Classification of Inkjet and Photographic Media. In *Proceedings of the International Conference on Digital Printing Technologies*, 2013.

L.J.P. van der Maaten. Barnes-Hut-SNE. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2013. [Oral acceptance rate: 21%; 275 citations]

L. Zhang and L.J.P. van der Maaten. Structure Preserving Object Tracking. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1838–1845, 2013. [Oral acceptance rate: 3.2%; 256 citations]

L.J.P. van der Maaten, M. Chen, S. Tyree, and K.Q. Weinberger. Learning with Marginalized Corrupted Features. In *Proceedings of the International Conference on Machine Learning (ICML), JMLR W&CP 28:410–418*, 2013. [Acceptance rate: 24%; 202 citations]

L.J.P. van der Maaten and K.Q. Weinberger. Stochastic Triplet Embedding. In *Proceedings of the IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, 2012. [255 citations]

J. Fang, A.L. Varbanescu, J. Shen, H. Sips, G. Saygili, and L.J.P. van der Maaten. Accelerating Cost Aggregation for Real-Time Stereo Matching. In *Proceedings of the IEEE International Conference on Parallel and Distributed Systems (ICPADS)*, pages 472–481, 2012.

L.J.P. van der Maaten. Audio-Visual Emotion Challenge 2012: A Simple Approach. In

*Proceedings of the International Conference on Multimodal Interaction (ICMI)*, pages 473–476, 2012.

J.M. Lewis, L.J.P. van der Maaten, and V.R. de Sa. A Behavioral Investigation of Dimensionality Reduction. In *Proceedings of the Cognitive Science Society (CSS)*, pages 671-676, 2012. [Oral acceptance rate: 32%; 50 citations]

G. Saygili, L.J.P. van der Maaten, and E.A. Hendriks. Improving Segment-based Stereo Matching using SURF Key Points. In *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, 2012.

G. Saygili, L.J.P. van der Maaten, and E.A. Hendriks. Feature-Based Stereo Matching Using Graph Cuts. In *Proceedings of the ASCI Conference, 2011*.

L. Zhang, L.J.P. van der Maaten, N. Koenderink, F. Golbach and E.A. Hendriks. Pose Estimation and Tracking of Eating Persons in Real-life Settings. In *Proceedings of the ASCI Conference, 2011*.

L.J.P. van der Maaten. Learning Discriminative Fisher Kernels. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 217–224, 2011. [Acceptance rate: 25%; 98 citations]

L.J.P. van der Maaten, M. Welling, and L.K. Saul. Hidden-Unit Conditional Random Fields. In *Proceedings of the International Conference on Artificial Intelligence & Statistics (AI-STATS)*, JMLR W&CP 15:479–488, 2011. [Acceptance rate: 28%; 98 citations]

L.J.P. van der Maaten. Discussion of “Spectral Dimensionality Reduction via Maximum Entropy” (*invited paper*). In *Proceedings of the International Conference on Artificial Intelligence & Statistics (AI-STATS)*, JMLR W&CP 15:60–62, 2011. [Acceptance rate: 28%]

D. Hu, L.J.P. van der Maaten, Y. Cho, L.K. Saul, and S. Lerner. Latent Variable Models for Predicting File Dependencies in Large-Scale Software Development. In *Advances of Neural Information Processing Systems (NIPS)*, volume 23, pages 865–873, 2010. [Acceptance rate: 24%]

A. Gelfand, L.J.P. van der Maaten, Y. Chen, and M. Welling. On Herding and the Perceptron Cycling Theorem. In *Advances of Neural Information Processing Systems (NIPS)*, volume 23, pages 694–702, 2010. [Acceptance rate: 24%; 25 citations]

L.J.P. van der Maaten and E.O. Postma. Texton-Based Analysis of Paintings. In *SPIE Optical Engineering and Applications*, volume 7798-16, 2010. [19 citations]

R. Min, L.J.P. van der Maaten, Z. Yuan, A. Bonner, and Z. Zhang. Deep Supervised t-Distributed Embedding. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 791–798, 2010. [Acceptance rate: 26%; 65 citations]

L.J.P. van der Maaten and E.A. Hendriks. Capturing Appearance Variation in Active Appearance Models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 34–41, 2010. [21 citations]

L.J.P. van der Maaten, A.G. Lange, and P.J. Boon. Visualization and Automatic Typology Construction of Pottery Profiles. In *Proceedings of the Conference on Computer*

*Applications and Quantitative Methods in Archaeology (CAA)*, 2009.

L.J.P. van der Maaten. Preserving Local Structure in Gaussian Process Latent Variable Models. In *Proceedings of BeNeLearn-09*, pages 81–88, 2009.

L.J.P. van der Maaten. Learning a Parametric Embedding by Preserving Local Structure. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AI-STATS)*, *JMLR: W&CP* 5, pp. 384–391, 2009. [673 citations]

L.J.P. van der Maaten and E.O. Postma. Texton-based Texture Classification. In *Proceedings of the Belgian-Dutch Artificial Intelligence Conference (BNAIC)*, pages 213–220, 2007.

S. Vanderlooy, L.J.P. van der Maaten, and I. Sprinkhuizen-Kuyper. Off-line Learning with Transductive Confidence Machines: An Empirical Evaluation. In *Proceedings of the International Conference on Machine Learning and Data Mining*, pages 310–323, 2007. [34 citations]

L.J.P. van der Maaten and P.J. Boon. COIN-O-MATIC: A Fast and Reliable System for Coin Classification. In *Proceedings of the MUSCLE Coin Workshop*, pages 7–17, 2006. [79 citations]

L.J.P. van der Maaten and E.O. Postma. Towards Automatic Coin Classification. In *Proceedings of the EVA-Vienna*, pages 19–26, 2006. [44 citations]

L.J.P. van der Maaten, P.J. Boon, J.J. Paijmans, A.G. Lange, and E.O. Postma. Computer Vision and Machine Learning for Archaeology. In *Proceedings of the Conference on Computer Applications and Quantitative Methods in Archaeology (CAA)*, pages 361–367, 2006.

L.J.P. van der Maaten and E.O. Postma. Improving Automatic Writer Identification. In *Proc. of the Belgian-Dutch AI Conference*, pages 260–266, 2005. [57 citations]

## JOURNAL PUBLICATIONS

C. Guo, A.Y. Hannun, B. Knott, L. van der Maaten, M. Tygert, and R. Zhu. Secure multiparty computations in floating-point arithmetic. In *Information and Inference: A Journal of the IMA*, 2021. [Impact factor: 1.36; 24 citations]

G. Huang, Z. Liu, G. Pleis, L.J.P. van der Maaten, and K.Q. Weinberger. Convolutional Networks with Dense Connectivity. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. [Impact factor: 17.730; 383 citations]

K. van Hecke, G. de Croon, L.J.P. van der Maaten, D. Hennes, and D. Izzo. Persistent Self-Supervised Learning Principle: From Stereo to Monocular Vision for Obstacle Avoidance. In *International Journal of Micro Air Vehicles* 10(2):186–206, 2018. [Impact factor: 1.049]

W. Pei, H. Dibeklioglu, D.M.J. Tax, and L.J.P. van der Maaten. Multivariate Time-Series Classification Using the Hidden-Unit Logistic Model. *IEEE Transactions on Neural Networks and Learning Systems*, 2017. [Impact factor: 4.370; 85 citations]

N. Pezzotti, B.P.F. Lelieveldt, L.J.P. van der Maaten, T. Höllt, E. Eisemann, A. Vilanova. Approximated and User Steerable tSNE for Progressive Visual Analytics. *IEEE Transactions on Visualization and Computer Graphics* 23(7), 2017. [Impact factor: 1.4; 298 citations]

- W.M. Kouw, J.H. Krijthe, M. Loog, and L.J.P. van der Maaten. Feature-Level Domain Adaptation. *Journal of Machine Learning Research* 17(171):1–32, 2016. [Impact factor: 3.42; 71 citations]
- Y. Cheng, M.T. Wong, L.J.P. van der Maaten, and E.W. Newell. Categorical analysis of human T cell heterogeneity with one-dimensional soli-expression by nonlinear stochastic embedding. *Journal of Immunology* 196(2):924-932, 2016. [Impact factor: 4.922; 69 citations]
- L.J.P. van der Maaten and R.G. Erdmann. Automatic Thread-Level Canvas Analysis. *IEEE Signal Processing Magazine* 32(4):38-45, 2015. [Impact factor: 5.852; 30 citations]
- M. van Sebille, L.J.P. van der Maaten, L. Xie, K. Jarolimek, R. Santbergen, R.A.C.M.M. van Swaaij, K. Leifer, and M. Zemana. Nanocrystal Size Distribution Analysis from Transmission Electron Microscopy Images. *Nanoscale* 7(48):20593-20606, 2015. [Impact factor: 7.394]
- B.M. Hoonhout, M. Radermacher, F. Baart, and L.J.P. van der Maaten. An Automated Method for Semantic Classification of Regions in Coastal Images. *Coastal Engineering* 105:1-12, 2015. [Impact factor: 2.428; 50 citations]
- G. Saygili, L.J.P. van der Maaten, and E.A. Hendriks. Adaptive Stereo Similarity Fusion using Confidence Measures. *Computer Vision and Image Understanding* 135:95-108, 2015. [Impact factor: 1.54; 23 citations]
- C.C. Laczny, T. Sternal, V. Plugaru, P. Gawron, A. Atashpendar, H.H. Margossian, S. Coronado, L. van der Maaten, N. Vlassis, and P. Wilmes. VizBin – An Application for Reference-Independent Visualization and Human-Augmented Binning of Metagenomic Data. *Microbiome* 3(1), 2015. [210 citations]
- A. Mahfouz, M. van de Giessen, L.J.P. van der Maaten, S. Huisman, M.J.T. Reinders, M.J. Hawrylycz, and B.P.F. Lelieveldt. Visualizing the spatial gene expression organization in the brain through non-linear similarity embeddings. *Methods* 73:79–89, 2015. [Impact factor: 4.197; 74 citations]
- L.J.P. van der Maaten. Accelerating t-SNE using Tree-Based Algorithms. *Journal of Machine Learning Research* 15(Oct):3221–3245, 2014. [Impact factor: 3.42; 2,740 citations]
- L. Zhang and L.J.P. van der Maaten. Preserving Structure in Model-Free Tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36(4):756-769, 2014. [Impact factor: 4.795; 177 citations]
- W. Abdelmoula, K. Skraskova, B. Balluff, R. Carreira, E. Tolner, B.P.F. Lelieveldt, L.J.P. van der Maaten, H. Morreau, A. van den Maagdenberg, R. Heeren, L. McDonnell, and J. Dijkstra. Automatic Generic Registration of Mass Spectrometry Imaging Data to Histology using Nonlinear Stochastic Embedding. In *Analytical Chemistry* 86(18):9204–9211, 2014. [Impact factor: 5.825; 73 citations]
- C.R. Johnson, Jr., P. Messier, W.A. Sethares, A.G. Klein, C. Brown, A. Hoang Do, P. Klausmeyer, P. Abry, S. Jaffard, H. Wendt, S. Roux, N. Pustelnik, N. van Noord, L.J.P. van der Maaten, E.O. Postma, J. Coddington, L.A. Daffner, H. Murata, H. Wilhelm, S. Wood, and M. Messier. Pursuing Automated Classification of Historic Photographic Papers from Raking Light Images. In *Journal of the American Institute*

*for Conservation* 53(3):159–170, 2014. [46 citations]

M. Mehu and L.J.P. van der Maaten. Multimodal Integration of Audio-Visual Cues in the Communication of Agreement and Disagreement. In *Journal of Nonverbal Behavior* 38(4):569–597, 2014. [Impact factor: 1.344; 29 citations]

J.M. Lewis, V.R. de Sa, and L.J.P. van der Maaten. Divvy: Fast and Intuitive Exploratory Data Analysis. *Journal of Machine Learning Research* 14(Oct):3159–5163, 2013. [Impact factor: 3.420]

M. Grimes, W.J. Lee, L.J.P. van der Maaten, and P. Shannon. Wrangling Phosphoproteomic Data to Elucidate Cancer Signaling Pathways. *PLoS One* 8(1):e52884, 2013. [Impact factor: 4.092; 18 citations]

L.J.P. van der Maaten and G.E. Hinton. Visualizing Non-Metric Similarities in Multiple Maps. *Machine Learning* 87(1):33–35, 2012. [Impact factor: 1.956; 300 citations]

L.J.P. van der Maaten and E.A. Hendriks. Action Unit Classification using Active Appearance Models and Conditional Random Fields. *Cognitive Processing* 13:507–518, 2012. [Impact factor: 1.03; 46 citations]

L.J.P. van der Maaten, M. Mahecha, and S. Schmidlein. Analyzing Floristic Inventories with Multiple Maps. *Ecological Informatics* 9:1–10, 2012. [Impact factor: 1.432]

O. Brinkkemper, L.J.P. van der Maaten, and P.J. Boon. Identification of Myosotis seeds by means of digital image analysis. *Vegetation History and Archaeobotany* 20(5):435–445, 2011. [Impact factor: 1.656]

P.J. Boon, A.G. Lange, L.J.P. van der Maaten, J.J. Paijmans, and E.O. Postma. Digital Support for Archaeology. *Interdisciplinary Science Review* 34(2-3):189–205, 2009. [Impact factor: 0.36]

L.J.P. van der Maaten and G.E. Hinton. Visualizing High-Dimensional Data Using t-SNE. *Journal of Machine Learning Research* 9(Nov):2579–2605, 2008. [Impact factor: 2.949; 38,711 citations]

## THESES

L.J.P. van der Maaten. *Feature Extraction from Visual Data*. Ph.D. thesis (*cum laude*), Tilburg University, The Netherlands, June 23<sup>rd</sup> 2009.

L.J.P. van der Maaten. *Improving Automatic Writer Identification*. M.Sc. thesis, Maastricht University, The Netherlands, August 13<sup>th</sup> 2005.

## SELECTED TECHNICAL REPORTS

L.J.P. van der Maaten, E.O. Postma, and H.J. van den Herik. Dimensionality Reduction: A Comparative Review. Tilburg University Technical Report, TICC-TR 2009-005, 2009. [3,645 citations]

L.J.P. van der Maaten. An Introduction to Dimensionality Reduction Using Matlab. Technical Report MICC 07–07. Maastricht University, Maastricht, The Netherlands, 2007. [199 citations]

## INVITED TALKS

Meta Llama 3. Air Street NYC AI Meetup, May 8<sup>th</sup> 2024, New York NY.

From Llama to Emu: Developing Generative AI at Meta. Machines Can See, April



17<sup>th</sup> 2024, Dubai, UAE.

Privacy-Preserving Machine Learning with CrypTen and Fisher Information Leakage. September 14<sup>th</sup> 2023, Cornell University, Ithaca NY.

Large-Scale Weakly Supervised Pretraining of Image Recognition Models. February 5<sup>th</sup>, 2023. Indian Statistical Institute, Kolkata, India.

Theoretically Speaking Panel with Sebastian Bubeck and Melanie Mitchell. July 22<sup>nd</sup> 2022, Simons Institute, Berkeley CA.

AI and Visualization Panel with Elijah Meeks, Octavia Camps, and Daniel Haehn. July 25<sup>th</sup> 2022, Northeastern University, Portland ME.

A Systematic Study of Bias Amplification. June 20<sup>th</sup> 2022, CVPR Workshop on Fairness and Trustworthiness in Computer Vision.

Large-Scale Clustering and Visualization. July 7<sup>th</sup>, 2021. European Conference on Data Analysis.

Fixes That Fail: Self-Defeating Improvements in Machine-Learning Systems. October 11<sup>th</sup> 2021, ICCV Tutorial on Cross-Model Compatibility in Computer Vision.

Computer Vision and Machine Learning: Potential Impact on Cultural Heritage. December 2<sup>nd</sup> 2021, Collection Care, Valencia, Spain.

On Private Prediction and Certified Removal. February 8<sup>th</sup> 2021. AAAI 2021 Workshop: Towards Robust, Secure and Efficient Machine Learning.

CrypTen. December 6<sup>th</sup>, 2020. NeurIPS 2020 Expo.

Privacy and Security in Machine Learning. December 3<sup>rd</sup>, 2020. AI Journey, Moscow, Russia.

Exploring the Limits of Weakly Supervised Learning. October 21<sup>st</sup>, 2020. IEEE/IBM What's Next in AI Symposium.

Using t-SNE for Data Analysis: A Quick Introduction. September 30<sup>th</sup>, 2020. BioTuring Webinar.

The Trade-Offs of Private Prediction. September 26<sup>th</sup>, 2020. OpenMined Privacy Conference.

CrypTen. July 30<sup>th</sup>, 2020. Responsible Data Summit.

From Visual Recognition to Visual Reasoning. July 1<sup>st</sup>, 2020. Machines Can See, Moscow, Russia.

From Visual Recognition to Visual Reasoning. July 1<sup>st</sup>, 2020. Machine Learning Summer School, Bilkent University, Turkey.

Adversarial Robustness: The End of the Early Years. June 19<sup>th</sup>, 2020. CVPR 2020 Workshop on Adversarial Attacks.

From Visual Recognition to Visual Reasoning. November 18<sup>th</sup>, 2019. Duke University, Durham NC.

Privacy-Preserving Machine Learning with CryptTen. October 10<sup>th</sup>, 2019. PyTorch DevCon, San Francisco CA.

From Visual Recognition to Visual Reasoning. September 11<sup>th</sup>, 2019. German Conference on Pattern Recognition (GCPR), Dortmund, Germany.

From Visual Recognition to Visual Reasoning. September 9<sup>th</sup>, 2019. Tilburg University, The Netherlands.

Does Object Recognition Work for Everyone? June 16<sup>th</sup>, 2019. CVPR Workshop on Global Challenges for Computer Vision, Long Beach CA.

Exploring the Limits of Weakly Supervised Learning. June 16<sup>th</sup>, 2019. CVPR Workshop on Visual Understanding by Learning from Web Data, Long Beach CA.

Making Image Recognition Work in the Real World. June 16<sup>th</sup>, 2019. CVPR DeepVision Workshop, Long Beach CA.

Developing Efficient Convolutional Networks (And Training Them at Scale). May 24<sup>th</sup>, 2019. HHMI Janelia Farm, Ashburn VA.

Developing Efficient Convolutional Networks (And Training Them at Scale). May 21<sup>st</sup>, 2019. Mount Sinai Hospital, New York NY.

Developing Efficient Convolutional Networks (And Training Them at Scale). April 5<sup>th</sup>, 2019. University of Amsterdam, The Netherlands.

From Visual Recognition to Visual Reasoning. April 4<sup>th</sup>, 2019. Erasmus University Rotterdam, The Netherlands.

Binary Image Selection: Interpretable Evaluation of Visual Grounding. February 25<sup>th</sup>, 2019. CVPR AC Meeting, San Diego CA.

From Visual Recognition to Visual Reasoning. November 3<sup>rd</sup>, 2018. Machine Intelligence Conference, MIT Media Lab, Cambridge MA.

Exploring the Limits of Convolutional Networks. August 30<sup>th</sup>, 2019. Deep Learning Sydney, Australia.

Visualizing Data using t-SNE. August 1<sup>st</sup>, 2018. Joint Statistical Meetings, Vancouver, Canada.

How Modern Computer Vision and Machine Learning May Impact the Analysis of Art. July 22<sup>th</sup>, 2018. Gordon Research Conference, Barcelona, Spain.

Do's and Don'ts of using t-SNE to Understand Vision Models. June 18<sup>th</sup>, 2018. CVPR Workshop on Interpretable Machine Learning for Computer Vision, Salt Lake City UT.

Image Classification: Models and Data. June 13<sup>th</sup>, 2018. Delft University of Technology, The Netherlands.

Computer Vision for the Arts. April 12<sup>th</sup>, 2018. The Frick Collection, New York NY.

From Visual Recognition to Visual Reasoning. April 6<sup>th</sup>, 2018. Awesome IT Symposium, Amsterdam, The Netherlands.

From Visual Recognition to Visual Reasoning. March 5<sup>th</sup>, 2018. CILVR Seminar, New York University, New York NY.

Two Tales about Image Classification. November 13<sup>th</sup>, 2017. Vision and Autonomous Systems Seminar, Carnegie-Mellon University, Pittsburgh PA.

From Visual Recognition to Visual Understanding. November 9<sup>th</sup>, 2017. Belgian-Dutch Artificial Intelligence Conference (keynote), Groningen, The Netherlands.

Two Tales about Image Classification. November 2<sup>nd</sup>, 2017. Bay Area Vision Meeting (BAVM), Menlo Park CA.

Efficient Recognition of Large Numbers of Visual Classes. October 13<sup>th</sup>, 2017. Machine Learning Meetup, Amsterdam, The Netherlands.

From Visual Recognition to Visual Understanding. October 11<sup>th</sup>, 2017. World Summit AI, Amsterdam, The Netherlands.

Building Image-Classification Models that are Accurate and Efficient. June 27<sup>rd</sup>, 2017. IBM T.J. Watson Research Center, Yorktown Heights NY.

Building Image-Classification Models that are Accurate and Efficient. June 23<sup>rd</sup>, 2017. ACM MobiSys (keynote), Niagara Falls NY.

The Trade-Offs of Deep-Learning Frameworks. June 5<sup>th</sup>, Semiconductor Research Company, University of Texas, Dallas TX.

Building Image-Classification Models that are Accurate and Efficient. April 28<sup>th</sup>, 2017. Big Data Meetup, New York NY.

Visualizing Data using Embeddings. March 28<sup>th</sup>, 2017. Applied Probability and Statistics Colloquium, Royal Holloway University, Egham, United Kingdom.

Understanding and Improving Convolutional Networks. March 2<sup>nd</sup>, 2017, Data Science Seminar, Leiden University, The Netherlands.

Visualizing Data using Embeddings. March 1<sup>st</sup>, 2017, Machine Learning Seminar, Linköping University, Sweden.

Visualizing Data using Embeddings. February 15<sup>th</sup>, 2017, Applied Math Seminar, Yale University, New Haven CT.

Learning to Solve Vision without Annotating Millions of Images. October 21<sup>st</sup>, 2016, Cornell University, Ithaca NY.

Learning to Solve Vision without Annotating Millions of Images. October 10<sup>th</sup>, 2016, Fourth ECCV Workshop on Web-Scale Vision and Social Media (invited keynote talk), Amsterdam, The Netherlands.

Weakly Supervised Training of Convolutional Networks. April 25<sup>th</sup>, 2016, Princeton University, Princeton NJ.

Visualizing Data using Embeddings. April 12<sup>th</sup>, 2016, Columbia University, New York NY.

Learning to Solve Vision without Annotating Millions of Images. March 22<sup>nd</sup>, 2016, ICT-Open, Amersfoort, The Netherlands.

Deep Learning: Embeddings. February 23<sup>rd</sup>, 2016, New York University, New York NY.

Visualizing Data using Embeddings. Design@Large Talks. February 10<sup>th</sup>, 2016, University of California, San Diego CA.

News Feed Ranking: Large-Scale Logistic Regression at Facebook. September 8<sup>th</sup>, 2015, Cornell Tech, New York NY.

Learning Embeddings for Data Visualization and Language Modeling. January 28<sup>th</sup>, 2015, Netherlands Institute for Neuroscience, Amsterdam, The Netherlands.

Constructing Maps to Visualize Big Data. January 15<sup>th</sup>, 2015, National Conference on the Mathematics of Operations Research, Lunteren, The Netherlands.

New Technology to Look Inside Paintings. September 25<sup>th</sup> 2014, eCult Dialogue Day, Dubrovnik, Croatia.

Visualizing Spatial Gene Expression Organization in the Brain with t-SNE. September 23<sup>rd</sup> 2014, Delft Data Science Symposium on Health, The Netherlands.

Visualizing (Big) Data using t-SNE. June 2<sup>nd</sup> 2014, Royal Netherlands Academy of Science (KNAW), Amsterdam, The Netherlands.

Learning with Marginalized Corrupted Features. May 29<sup>th</sup> 2014, Deepmind, London, United Kingdom.

t-SNE and MCF: Constructing and Corrupting Features. March 31<sup>st</sup> 2014, Facebook, Menlo Park CA.

Visualizing Data using t-SNE. January 13<sup>th</sup> 2014, Leiden University Medical Center, The Netherlands.

Looking at Paintings through the Eyes of a Computer. October 25<sup>th</sup> 2013, University of Amsterdam, The Netherlands.

Visualizing Data using t-SNE. October 17<sup>th</sup> 2013, Washington University, St. Louis MO.

Visualizing Data using t-SNE. October 15<sup>th</sup> 2013, Mind Research Network / University of New Mexico, Albuquerque NM.

Learning with Marginalized Corrupted Features. July 10<sup>th</sup> 2013, ROKS Workshop, Leuven, Belgium.

Visualizing Data using t-SNE. June 24<sup>th</sup> 2013, Google Tech Talk, Mountain View CA. [The video of this talk has been viewed over 140,000 times on Youtube.]

Structure-Preserving Object Tracking and Marginalized Corrupted Features. June 10<sup>th</sup> 2013, INRIA-LEAR, Grenoble, France.

New Developments in Tracking and Forensic Painting Analysis. June 6<sup>th</sup> 2013, Netherlands Forensic Institute and Dutch National Police, The Hague, The Netherlands.

Learning with Marginalized Corrupted Features. May 30<sup>th</sup> 2013, Dutch Association for Pattern Recognition and Image Processing, Utrecht, The Netherlands.

Looking at Paintings through the Eyes of a Computer. April 25<sup>th</sup> 2013, Netherlands Institute for Advanced Study, Wassenaar, The Netherlands.

Learning with Marginalized Corrupted Features. April 18<sup>th</sup> 2013, Centrum voor Wiskunde en Informatica (CWI), Amsterdam, The Netherlands.

On Constructing and Corrupting Features. April 1<sup>st</sup> 2013, Toyota Technological Institute, Chicago IL.

How to Visualize High-Dimensional Data? January 24<sup>th</sup> 2013, Data Science Meetup, Utrecht, The Netherlands.

Learning with Marginalized Corrupted Features. November 9<sup>th</sup> 2012, Radboud University Nijmegen, The Netherlands.

Machine Learning for Data Visualization: New Techniques and Applications. October 26<sup>th</sup> 2012, SNN Machine Learning Award Lecture, Maastricht, The Netherlands.

Learning by Marginalizing Corrupted Features. October 20<sup>th</sup> 2012, University of California, San Diego.

Learning when Sensors Break Down. June 22<sup>nd</sup> 2012, European Space Association (ESA-ESTEC), Noordwijk, The Netherlands.

Stochastic Triplet Embedding. April 20<sup>th</sup> 2012, Dutch Association for Pattern Recognition and Image Processing, Rijswijk, The Netherlands.

Learning Embeddings for Data Visualization. March 27<sup>th</sup> 2012, University of Amsterdam, The Netherlands.

Introduction to Embedding for Visualization. February 20<sup>th</sup> 2012, Dagstuhl Seminar 12081, Germany.

Hidden-Unit Conditional Random Fields. January 27<sup>th</sup> 2012, University of Geneva, Switzerland.

Visualizing Data using t-SNE. October 17<sup>th</sup> 2011, University of Groningen, The Netherlands.

Tutorial on Machine Learning for Data Visualization. August 22<sup>nd</sup> 2011, Dagstuhl Seminar 11341, Germany.

Conditional Random Fields with Hidden Units. June 15<sup>th</sup> 2011, Imperial College London, United Kingdom.

Hidden-Unit Conditional Random Fields. May 18<sup>th</sup> 2011, Dutch Association for Pattern Recognition and Image Processing, Utrecht, The Netherlands.

Computer Techniques for Analyzing Archaeological Finds. February 17<sup>th</sup> 2011, University of Amsterdam, The Netherlands.

Conditional Random Fields with Hidden Units. December 23<sup>rd</sup> 2010, Tilburg University, The Netherlands.

Conditional Random Fields with Hidden Units. November 22<sup>nd</sup> 2010, University of California, San Diego.

Visualizing Data using t-SNE. November 8<sup>th</sup> 2010, University of California, Irvine.

Visualizing Data using t-SNE. April 30<sup>th</sup> 2010, University of Southern California, Los Angeles.

Exploring High-Dimensional Data using Dimensionality Reduction. June 9<sup>th</sup> 2009, Workshop on Novel Data Mining Strategies for Biochemical Cycles and Biosphere-Atmosphere Interactions, Jena, Germany.

Visualizing Data using t-SNE. March 13<sup>th</sup> 2009, Radboud University Nijmegen, The Netherlands.

On Parametric and Non-Metric Variants of t-SNE. April 6<sup>th</sup> 2009, University of Toronto.

Visualizing Data using t-SNE. March 30<sup>th</sup> 2009, University of California, San Diego.

Texton-based analysis and visualization of Van Gogh's paintings. October 21<sup>st</sup> 2008, Van Gogh Museum, Amsterdam, The Netherlands.

**REFERENCES** Available on request.