







Ishan Misra

Research Scientist @ GenAI (Meta)

I am a Research Scientist at Meta where I work on Computer Vision and Machine Learning. My research interest is in training multimodal machine learning models at scale.

 ishanmisra.github.io
 ishanmisra@gmail.com
 ishanmisra@meta.com
 [CV PDF](#)
 [LinkedIn](#)

 [@ishanmisra_](#)
 [@ishanmisra](#)
 [Google Scholar](#)

Education

Ph.D. in Robotics

Carnegie Mellon University, Pittsburgh, USA

Thesis: *Visual Learning with Minimal Human Supervision*

Advisor: Martial Hebert , Co-advisor: Abhinav Gupta

Committee: Martial Hebert , Abhinav Gupta , Deva Ramanan , Alyosha Efros , Andrew Zisserman

 [CMU School of Computing Science Distinguished Dissertation Award \(Runner Up\)](#)

 [Thesis Project Page](#)  [Thesis](#)

M.S. in Robotics

Carnegie Mellon University, Pittsburgh, USA

Thesis: *Data Driven Exemplar Model Selection*

Advisor: Martial Hebert

 [Siebel Scholarship](#)

BTech (Hons) in Computer Science Engineering

IIIT, India

Thesis: *Hybrid implementation of Image Dithering*

Advisor: P J Narayanan

Rank 1/150 in graduating class

 [Gold Medalist for Computer Science \(Summa Cum Laude\)](#)

Selected Honors and Awards

- 2022 MIT Tech Review's 35 innovators under 35
Contributions to self-supervised learning in computer vision. Compiled globally across the world in all technology areas.
- 2021 Best Paper Finalist CVPR 2021
Paper: "Audio Visual Instance Discrimination"
- 2018 Carnegie Mellon (SCS) Distinguished Dissertation, Runner Up
PhD Thesis: "Visual Learning with Minimal Human Supervision"
- 2014 Siebel Scholarship
Selected from all graduate students at the School of Computer Science, CMU
- 2014 Best Student Paper, IEEE WACV
Paper: Data Driven Exemplar Model Selection

Press

- November 2023 Emu Video: Text-to-Video Generation
Meta (announced by Mark Zuckerberg), Reuters, TheVerge, Tech Crunch
- April 2023 DINOv2: State-of-the-art computer vision models with self-supervised learning
Meta (announced by Mark Zuckerberg), Venture Beat, KD Nuggets, AI Papers Academy
- May 2023 ImageBind: a new way to link AI across senses
Meta (announced by Mark Zuckerberg), The Verge, Engadget, WorldOfAI
- March 2021 Self-supervised Learning: The dark matter of intelligence
Meta (co-written with Yann LeCun)
- March 2021 SEER: The start of a more powerful, flexible, and accessible era for computer vision
Meta, CNBC, Wired, Engadget

Industry Research Experience

- Feb 2023 - Present **Meta (Facebook)**, Seattle, WA, USA
Research Scientist, Generative AI
Lead the research team for video generation and understanding
- Sep 2018 - Feb 2023 **Meta (Facebook)**, New York, NY, USA
Research Scientist, FAIR

Fundamental research on self-supervised and multimodal representation learning

Summer 2017

Meta (Facebook), New York, NY, USA

Research Intern, FAIR

Mentor: Rob Fergus , Ross Girshick , Laurens van der Maaten

Designed a visual question answering system that acquires knowledge by identifying gaps in its knowledge and asking questions to humans to fill those gaps, published at CVPR'17

Summer 2015

Microsoft Research, Redmond, WA, USA

Research Intern, Computer Vision

Mentor: Ross Girshick , Larry Zitnick , Margaret Mitchell , Jacob Devlin

Worked on vision-language systems for generating questions (published at ACL), storytelling datasets (published at NAACL), weakly supervised learning (published at CVPR)

Summer 2014

Microsoft Research, Redmond, WA, USA

Research Intern, Computer Vision

Mentor: Xia Sheng Hua

Worked on improving image search systems using patch-based deep learning features (2 US patents granted)

Academic Research Experience

Summer 2012

INRIA, Paris, France

Research Intern, Ecole Centrale

Mentor: Iasonas Kokkinos

Worked on shape from shading

Summer 2011

Yale, New Haven, CT, USA

Research Intern, Computer Science

Mentor: Bryan Ford

Worked on deterministic distributed operating systems: writing bootloaders in assembly and parallel threading libraries.

Talks

Beyond pretty pictures: What's needed to make Generative Visual Models useful?

July 2024

ELLIS Workshop on Open Problems in Computer Vision & Generative Modelling at Munich, Germany

Generative models for Computer Vision

July 2024

Oxford Machine Learning Summer School at the University of Oxford, UK

- Improving generative models for vision: high quality videos and precise image control**
 March 2024 ELLIS Winter School on Foundation Models at Amsterdam, Netherlands
- Generative Models for Multimodal Learning**
 February 2024 Human-centric representation learning workshop at AAIL
- Emu Video: State of the Art Video Generation**
 January 2024 Aleksa Gordic's Discord (AI Epiphany)
- Using unlabeled data to scale representations across modalities**
 October 2023 Learning from Noisy and Unlabeled Data, ICCV, Paris, France
- SSL scales multimodal pretraining to more modalities and data**
 October 2023 Big Model Adaptation for Computer Vision Workshop, ICCV, Paris, France
 June 2023 Transformers for Vision Workshop, CVPR, Vancouver, Canada
- Machine Learning without Human Supervision**
 January 2023 Epoch Foundation, Young Innovators Conference
- General-purpose Visual Recognition Systems: Beyond a Single Modality and a Task**
 October 2022 CV in the Wild Workshop. ECCV, Tel Aviv, Israel
- General-purpose Visual Recognition Across Modalities with Limited Supervision**
 October 2022 L2ID Workshop. ECCV, Tel Aviv, Israel
- Representation Learning Beyond a Single Dataset and Modality**
 October 2022 Self-supervised Learning: What's Next Workshop. ECCV, Tel Aviv, Israel
- Self-supervised Learning**
 August 2022 Summer School at IIIT-H, India
- Visual Recognition and Self-supervised Learning**
 August 2022 Oxford Machine Learning Summer School, UK
- General purpose Vision Models**
 August 2022 Visual Geometry Group, University of Oxford
- Self-supervised Visual Learning**
 April 2022 Keynote speaker at Ghostday ML Conference
- Object Discovery using Transformers**
 2022 Invited at Zipline, Inc.

Can Machines Learn to See without Human Supervision?

2022 IIM Ahmedabad, India

3D Recognition using Transformers

2021 Invited at Aurora, Inc.

Redundancy Reduction for Self-supervised Learning

2021 University of Illinois, Urbana Champaign

Learning Vision Models with Minimal Supervision

2021 IIT Jodhpur, India

Guest Lectures

Video Diffusion Models

2023 Guest Lecture: Deep Learning course by Abhinav Gupta at Carnegie Mellon University

Multimodal learning

2023 Guest Lecture: Deep Learning course by Abhinav Shrivastava at UMD, College Park

Self-supervised Learning in Vision

2023, 2022, 2021, 2020, 2019 Guest Lecture: Deep Learning course by Rob Fergus at NYU

2021, 2020 Guest Lecture: Deep Learning course by Yann LeCun at NYU

2022, 2021 Guest Lecture: Deep Learning course by Zsolt Kira at Georgia Tech

2022, 2020, 2019 Guest Lecture: Deep Learning course by Dhruv Batra at Georgia Tech

2021, 2020, 2019 Guest Lecture: Deep Learning course by Abhinav Shrivastava at UMD, College Park

Structure from Motion

2019 Guest Lecture: Computer Vision course by Rob Fergus at NYU

Collaborators and Interns

January 2023 - present

Saketh Rambhatla at GenAI (Meta)

PhD from University of Maryland, College Park

Postdoctoral Researcher

Summer 2023

Xudong Wang at GenAI (Meta)

PhD, University of California, Berkeley

Internship on instance conditioned diffusion models. Co-hosted with Rohit Girdhar, Saketh Rambhatla.

- Summer 2023 **Sachit Menon** at GenAI (Meta)
PhD, University of California, Berkeley
Internship on LLMs + diffusion models for visual instruction generation. Co-hosted with Rohit Girdhar.
- Summer 2023 **Feng (Jeff) Liang** at GenAI (Meta)
PhD, University of Texas, Austin
Internship on video diffusion models for editing. Co-hosted with Bichen Wu.
- 2022-2023 **Nur Muhammad Mahi Shafiullah** at GenAI (Meta)
PhD, New York University
Visiting Researcher working on Home Robotics. Co-hosted with Soumith Chintala.
👤 Apple PhD Fellowship 2023
- Summer 2023 **Basile Van Hoorick** at GenAI (Meta)
PhD, Columbia University
Internship on video diffusion models. Co-hosted with Laurens van der Maaten.
- Summer 2022 **Xudong Wang** at FAIR (Meta)
PhD, University of California, Berkeley
Internship on self-supervised segmentation (published at CVPR'23). Co-hosted with Rohit Girdhar.
- Summer 2022 **Yue Zhao** at FAIR (Meta)
PhD, University of Texas, Austin
Internship on LLMs + Video understanding (published at CVPR'23). Co-hosted with Rohit Girdhar.
👤 NVIDIA Graduate Fellowship 2024-2025
- Summer 2021 **Xingyi Zhou** at FAIR (Meta)
PhD, University of Texas, Austin
Internship on Open Vocabulary Object Detection (published at ECCV'22). Co-hosted with Armand Joulin, Rohit Girdhar.
👤 Facebook Fellowship 2021
Now: Research Scientist at Google Research
- Summer 2021 **Bowen Cheng** at FAIR (Meta)
PhD, University of Illinois at Urbana-Champaign
Internship on transformers for pixel segmentation (published at CVPR'22)
Now: ML Scientist at Tesla
- Summer 2021 **Karan Desai** at FAIR (Meta)
PhD, University of Michigan at Ann Arbor

Internship on using scribbles for segmenting objects (published at BMVC'22). Co-hosted with Laurens van der Maaten.

Summer 2020

Zaiwei Zhang at FAIR (Meta)

PhD, University of Texas, Austin

Internship on self-supervised learning using depth (published at ICCV'21).

Co-hosted with Armand Joulin, Rohit Girdhar.

Now: Researcher at Cruise

Summer 2020

Zhongzheng (Jason) Ren at FAIR (Meta)

PhD, University of Illinois at Urbana-Champaign

Internship on weakly supervised learning for 3D detection (published at CVPR'21). Co-hosted with Rohit Girdhar.

Now: Research Scientist at Apple

Summer 2020

Yuki Asano at FAIR (Meta)

PhD, University of Oxford

Internship on self-supervised learning for object discovery. Co-hosted with Armand Joulin, Piotr Bojanowski, Andrea Vedaldi.

Now: Assistant Professor at University of Amsterdam

Summer 2019

Pedro Morgado at FAIR (Meta)

PhD, University of California, San Diego

Internship on self-supervised learning for audiovisual learning (published at ECCV'20, best-paper finalist).

Now: Assistant Professor at University of Wisconsin-Madison (EECS)

Summer 2019

Huaizu Jiang at FAIR (Meta)

PhD, University of Massachusetts, Amherst

Internship on visual question answering (published at CVPR'20). Co-hosted with Xinlei Chen.

Now: Assistant Professor at Northeastern University

Summer 2019

Jyh-Jing Hwang at FAIR (Meta)

PhD, University of California, Berkeley

Internship on anytime video recognition. Co-hosted with Laurens van der Maaten.

Now: Research Scientist at Waymo Research

Summer 2019

Yan Wang at FAIR (Meta)

PhD, Cornell University

Internship on anytime image recognition. Co-hosted with Laurens van der Maaten.

Now: Research Scientist at Waymo Research

Summer 2019

Terrance DeVries at FAIR (Meta)

PhD, Cornell University

Internship on bias in object recognition systems (published at CVPRW'2019). Co-hosted with Laurens van der Maaten.

Now: Research Scientist at Luma AI

2018 **Nilesh Kulkarni** at Carnegie Mellon University

Masters in Robotics, Carnegie Mellon University

Worked on image augmentations for object detection (published at ICCV'2019)

Now: PhD Student at University of Michigan, Ann Arbor

2016 **Debidatta Dwibedi** at Carnegie Mellon University

Masters in Robotics, Carnegie Mellon University

Worked on image augmentations for object detection (published at ICCV'2017)

Now: Research Scientist at Google Brain

Academic Service

Area Chair

IEEE/CVF Conference on Computer Vision and Pattern Recognition **CVPR**: 2021 , 2023 , 2024

European Conference on Computer Vision **ECCV**: 2024

Neural Information Processing Systems **NeurIPS**: 2022 , 2023 , 2024

International Conference on Learning Representations **ICLR**: 2023

Reviewing

IEEE/CVF Conference on Computer Vision and Pattern Recognition **CVPR**: 2015 , 2016 , 2017 , 2018 , 2019 , 2020 , 2022

IEEE/CVF International Conference on Computer Vision **ICCV**: 2015 , 2017 , 2019 , 2021 , 2023

European Conference on Computer Vision **ECCV**: 2016 , 2018 , 2020 , 2022

Neural Information Processing Systems **NeurIPS**: 2018 , 2019 , 2020 , 2021

International Conference on Learning Representations **ICLR**: 2019 , 2020 , 2021

Workshops

Workshop on Self-Supervised Learning Theory and Practice : 2023 , 2022 , 2021 , 2020

Tutorial on Self-Supervised Learning : 2023
Extreme Scale Vision : 2019
Vision & Language StoryTelling Workshop : 2017

Doctoral Thesis Committee

2024 Xudong Wang (University of California Berkeley)
2024 Alexandre Devillers (INRIA)
2023 Nirat Saini (University of Maryland, College Park)

Publications

A complete and updated list of my publications is available on my Google Scholar profile.

The Llama 3 Herd of Models

The Llama3 Team (played role of a Core Contributor for video recognition)
arxiv (arxiv). 2024.

[PDF](#) [Code](#)

Emu Video: Factorizing Text-to-Video Generation by Explicit Image Conditioning

Rohit Girdhar^{*}, Mannat Singh^{*}, Andrew Brown^{*}, Quentin Duval^{*},
Samaneh Azadi^{*}, Sai Saketh Rambhatla, Akbar Shah, Xi Yin, Devi Parikh,
[Ishan Misra](#)^{*}

ECCV (ECCV). 2024.

[PDF](#) [Code](#) [Colab](#) [BibTeX](#) ^{*}Authors contributed equally

FlowVid: Taming Imperfect Optical Flows for Consistent Video-to-Video Synthesis

Feng Liang, Bichen Wu, Jialiang Wang, Licheng Yu, Kunpeng Li, Yanan Zhao, [Ishan Misra](#), Jia-Bin Huang, Peizhao Zhang, Peter Vajda, Diana Marculescu

CVPR (CVPR). 2024.

[PDF](#) [BibTeX](#) [Highlight](#)

InstanceDiffusion: Instance-level Control for Image Generation

Xudong Wang, Trevor Darrell, Sai Saketh Rambhatla, Rohit Girdhar, [Ishan Misra](#)

CVPR (CVPR). 2024.

[PDF](#) [Code](#) [BibTeX](#)

Generating Illustrated Instructions

Sachit Menon, [Ishan Misra](#), Rohit Girdhar

CVPR (CVPR). 2024.

[PDF](#) [BibTeX](#)

VideoCutLER: Surprisingly Simple Unsupervised Video Instance Segmentation

Xudong Wang , [Ishan Misra](#) , Ziyun Zheng , Rohit Girdhar , Trevor Darrell

CVPR (CVPR). 2024.

[PDF](#) [BibTeX](#)

The effectiveness of MAE pre-pretraining for billion-scale pretraining

Mannat Singh * , Quentin Duval * , Kalyan Vasudev Alwala * , Haoqi Fan , Vaibhav Aggarwal , Aaron Adcock , Armand Joulin , Piotr Dollár , Christoph Feichtenhofer , Ross Girshick , Rohit Girdhar , [Ishan Misra](#)

ICCV (ICCV). 2023.

[PDF](#) [Code](#) [Colab](#) [BibTeX](#) *Authors contributed equally

MOST: Multiple Object localization with Self-supervised Transformers for object discovery.

Sai Saketh Rambhatla , [Ishan Misra](#) , Rama Chellappa , Abhinav Shrivastava

ICCV (ICCV). 2023.

[PDF](#) [Code](#) [BibTeX](#) [Oral](#)

MonoNeRF: Learning Generalizable NeRFs from Monocular Videos without Camera Poses

Yang Fu , [Ishan Misra](#) , Xiaolong Wang

ICML (ICML). 2023.

[PDF](#) [BibTeX](#)

ImageBind: One Embedding Space To Bind Them All

Rohit Girdhar * , Alaaeldin El-Nouby * , Zhuang Liu , Mannat Singh , Kalyan Vasudev Alwala , Armand Joulin , [Ishan Misra](#) *

CVPR (CVPR). 2023.

[Demo](#) [PDF](#) [Code](#) [Demo](#) [BibTeX](#) [Highlighted paper](#) *Authors contributed equally

Cut and Learn for Unsupervised Object Detection and Instance Segmentation

Xudong Wang , Rohit Girdhar , Stella X. Yu , [Ishan Misra](#)

CVPR (CVPR). 2023.

[PDF](#) [Code](#) [BibTeX](#)

Learning Video Representations from Large Language Models

Yue Zhao , [Ishan Misra](#) , Philipp Krahenbuhl , Rohit Girdhar

CVPR (CVPR). 2023.

[PDF](#) [Code](#) [Colab](#) [BibTeX](#) [Highlighted paper](#)

The Hidden Uniform Cluster Prior in Self-Supervised Learning

Mahmoud Assran , Randall Balestriero , Quentin Duval , Florian Bordes ,
Ishan Misra , Piotr Bojanowski , Pascal Vincent , Michael Rabbat , Nicolas
Ballas

ICLR (ICLR). 2023.

 PDF  BibTeX

OmniMAE: Single Model Masked Pretraining on Images and Videos

Rohit Girdhar *, Alaaeldin El-Nouby *, Mannat Singh *, Kalyan Vasudev Alwala
*, Armand Joulin , Ishan Misra *

CVPR (CVPR). 2023.

 PDF  Code  BibTeX *Authors contributed equally

Masked Siamese Networks for Label-Efficient Learning

Mahmoud Assran , Mathilde Caron , Ishan Misra , Piotr Bojanowski , Florian
Bordes , Pascal Vincent , Armand Joulin , Michael Rabbat , Nicolas Ballas

ECCV (ECCV). 2022.

 PDF  Code  BibTeX

Detecting Twenty-thousand Classes using Image-level Supervision

Xingyi Zhou , Rohit Girdhar , Armand Joulin , Phillip Krahenbuhl , Ishan Misra

ECCV (ECCV). 2022.

 PDF  Code  BibTeX

Vision Models Are More Robust And Fair When Pretrained On Uncurated Images Without Supervision

Priya Goyal , Quentin Duval , Isaac Seessel , Mathilde Caron , Ishan Misra ,
Levent Sagun , Armand Joulin , Piotr Bojanowski

Arxiv (Arxiv). 2022.

 PDF

Omnivore: A Single Model for Many Visual Modalities

Rohit Girdhar *, Mannat Singh *, Nikhila Ravi *, Laurens van der Maaten ,
Armand Joulin , Ishan Misra *

CVPR (CVPR). 2022.

 PDF  Code  BibTeX  Oral *Authors contributed equally

Masked-attention Mask Transformer for Universal Image Segmentation

Bowen Cheng , Ishan Misra , Alexander G. Schwing , Alexander Kirillov ,
Rohit Girdhar

CVPR (CVPR). 2022.

 PDF  Code  BibTeX

An End-to-End Transformer Model for 3D Object Detection

Ishan Misra , Rohit Girdhar , Armand Joulin

ICCV (ICCV). 2021.

 PDF  Code  BibTeX  Oral

Emerging Properties in Self-Supervised Vision Transformers

Mathilde Caron , Hugo Touvron , [Ishan Misra](#) , Hervé Jégou , Julien Mairal , Piotr Bojanowski , Armand Joulin

ICCV (ICCV). 2021.

 PDF  Code

Self-Supervised Pretraining of 3D Features on any Point-Cloud

Zaiwei Zhang , Rohit Girdhar , Armand Joulin , [Ishan Misra](#)

ICCV (ICCV). 2021.

 PDF  Code  BibTeX

MDETR : Modulated Detection for End-to-End Multi-Modal Understanding

Aishwarya Kamath , Mannat Singh , Yann LeCun , [Ishan Misra](#) , Gabriel Synnaeve , Nicolas Carion

ICCV (ICCV). 2021.

 PDF  Code  Oral

Audio-Visual Instance Discrimination with Cross-Modal Agreement

Pedro Morgado , Nuno Vasconcelos , [Ishan Misra](#)

CVPR (CVPR). 2021.

 PDF  Code  BibTeX  Best Paper Candidate

Robust Audio-Visual Instance Discrimination

Pedro Morgado , [Ishan Misra](#) , Nuno Vasconcelos

CVPR (CVPR). 2021.

 PDF  BibTeX  Oral

Barlow Twins: Self-Supervised Learning via Redundancy Reduction

Jure Zbontar * , Li Jing * , [Ishan Misra](#) , Yann LeCun , Stéphane Deny

ICML (ICML). 2021.

 PDF  Code  BibTeX *Authors contributed equally

3D Spatial Recognition without Spatially Labeled 3D

Zhongzheng Ren , [Ishan Misra](#) , Alexander G. Schwing , Rohit Girdhar

CVPR (CVPR). 2021.

 PDF

Unsupervised Learning of Visual Features by Contrasting Cluster Assignments

Mathilde Caron , [Ishan Misra](#) , Julien Mairal , Priya Goyal , Piotr Bojanowski , Armand Joulin

NeurIPS (NeurIPS). 2020.

 PDF  Code  BibTeX

Self-Supervised Learning of Pretext-Invariant Representations

[Ishan Misra](#) , Laurens van der Maaten

CVPR (CVPR). 2020.

 PDF  Code  BibTeX

ClusterFit: Improving Generalization of Visual Representations

Xueting Yan *, [Ishan Misra](#) *, Abhinav Gupta , Deepti Ghadiyaram *, Dhruv Mahajan *

CVPR (CVPR). 2020.

 PDF  Code  BibTeX *Authors contributed equally

In Defense of Grid Features for Visual Question Answering

Huaizu Jiang , [Ishan Misra](#) , Marcus Rohrbach , Erik Learned-Miller , Xinlei Chen

CVPR (CVPR). 2020.

 PDF  Code  BibTeX

3D-RelNet: Joint Object and Relational Network for 3D Prediction

Nilesh Kulkarni , [Ishan Misra](#) , Shubham Tulsiani , Abhinav Gupta

ICCV (ICCV). 2019.

 PDF  Code  BibTeX

Scaling and Benchmarking Self-Supervised Visual Representation Learning

Priya Goyal , Dhruv Mahajan , Abhinav Gupta *, [Ishan Misra](#) *

ICCV (ICCV). 2019.

 PDF  Code  BibTeX *Authors contributed equally

Binary Image Selection (BISON): Interpretable Evaluation of Visual Grounding

Hexiang Hu , [Ishan Misra](#) , Laurens van der Maaten

ICCV Workshop on Vision and Language (ICCV Workshop on Vision and Language). 2019.

 PDF  Code  BibTeX

Does Object Recognition Work for Everyone?

Terrance DeVries *, [Ishan Misra](#) *, Changan Wang *, Laurens van der Maaten

CVPR (CVPR). 2019.

 PDF  BibTeX *Authors contributed equally

Mainstream: Dynamic Stem-Sharing for Multi-Tenant Video Processing

Angela Jiang , Daniel L.-K. Wong , Christopher Canel , [Ishan Misra](#) , Michael Kaminsky , Michael Kozuch , Padmanabhan Pillai , David G. Andersen and Gregory Ganger

USENIX Annual Technical Conference (USENIX Annual Technical Conference). 2018.

 PDF  BibTeX

Learning by Asking Questions

[Ishan Misra](#) , Ross Girshick , Rob Fergus , Martial Hebert , Abhinav Gupta , Laurens van der Maaten

CVPR (CVPR). 2018.

 PDF  BibTeX  Oral

Cut Paste and Learn: Surprisingly Easy Synthesis for Instance Detection

Debidatta Dwibedi , [Ishan Misra](#) , Martial Hebert

ICCV (ICCV). 2017.

 PDF  Code  BibTeX

From Red Wine to Red Tomato: Composition with Context

[Ishan Misra](#) , Abhinav Gupta , Martial Hebert

CVPR (CVPR). 2017.

 PDF  Code  BibTeX  Oral

Shuffle and Learn: Unsupervised Learning using Temporal Order Verification

[Ishan Misra](#) , C. Lawrence Zitnick and Martial Hebert

ECCV (ECCV). 2016.

 PDF  Code  BibTeX

Seeing through the Human Reporting Bias: Visual Classifiers from Noisy

[Ishan Misra](#) , C. Lawrence Zitnick , Margaret Mitchell and Ross Girshick

CVPR (CVPR). 2016.

 PDF  Code  BibTeX

Cross-stitch Networks for Multi-Task Learning

[Ishan Misra](#) , Abhinav Shrivastava , Abhinav Gupta and Martial Hebert

CVPR (CVPR). 2016.

 PDF  BibTeX  Spotlight

Generating Natural Questions About an Image

Nasrin Mostafazadeh , [Ishan Misra](#) , Jacob Devlin ,

ACL (ACL). 2016.

 PDF  Code  BibTeX  Oral Long Paper

Visual Storytelling

Ting-Hao Huang , Francis Ferraro , Nasrin Mostafazadeh , [Ishan Misra](#) ,

Jacob Devlin , Aishwarya Agrawal ,

Ross Girshick ,

NAACL (NAACL). 2016.

 PDF  BibTeX

Watch and Learn: Semi-Supervised Learning of Object Detectors from Video

[Ishan Misra](#) , Abhinav Shrivastava and Martial Hebert

CVPR (CVPR). 2015.

 PDF  BibTeX

Applying artificial vision models to human scene understanding

E. Aminoff , M. Toneva , A. Shrivastava , X. Chen , I. Misra ,

Journal of Frontiers in Computational Neuroscience (Journal of Frontiers in Computational Neuroscience). 2015.

 PDF

Data-driven Exemplar Model Selection

[Ishan Misra](#) , Abhinav Shrivastava and Martial Hebert

WACV (WACV). 2014.

 PDF  BibTeX  Best Student Paper