# Tanmay Gupta

**Work Email:** tanmayg@allenai.org ● **Website:** http://tanmaygupta.info

**CURRENT ROLE**

**Research Scientist, Allen Institute for Artificial Intelligence** Jul 2020 – Present
Researcher in the PRIOR team working on general-purpose learning systems for vision and language.

**EDUCATION**

**University of Illinois Urbana Champaign**, Illinois, USA
PhD in Computer Science Aug 2014 – May 2020
Adviser: Derek Hoiem
Thesis: Representations from Vision and Language
Thesis Committee: Derek Hoiem, Abhinav Gupta, Svetlana Lazebnik, Alexander Schwing

**Indian Institute of Technology Kanpur**, Uttar Pradesh, India
Bachelor of Technology (B.Tech.) in Electrical Engineering Jul 2010 – May 2014
Thesis: Face Tracking and Recognition with Pose and Illumination Variations
Thesis Advisor: Aditya K. Jagannatham
CPI (Cumulative Performance Index): 9.9 / 10.0

**RESEARCH INTERESTS**

**General Purpose Vision:** Design, creation, and evaluation of general-purpose vision and language systems that can learn a large number of skills and concepts efficiently and without forgetting and without requiring any change to the network architecture.

**INTERNSHIPS**

**Research Intern, Learning and Perception Research, Nvidia** May 2019 – Dec 2019
*Collaborators*: Arash Vahdat, Xiaodong Yang, Gal Chechik, Jan Kautz
- Developed a novel contrastive learning approach for weakly supervised phrase grounding.

**Research Intern, Allen Institute for Artificial Intelligence** May 2017 – Dec 2017
*Collaborators*: Aniruddha Kembhavi, Dustin Schwenk, Ali Farhadi, Derek Hoiem
- Proposed and implemented a semi-parametric approach for text to video generation. We demonstrated this approach for a new richly annotated FLINTSTONES dataset.

**Software Development Intern, Visual Search Team, A9.com** May 2015 – Aug 2015
*Manager*: C. J. Taylor, *Mentor*: Michael Lou
- Developed an algorithm for depth estimation and foreground segmentation from a video stream.

**Undergraduate Research Intern, Cornell University** May 2013 – Jul 2013
*Supervisor*: Prof. Tsuhan Chen, Professor and Director of ECE, Cornell University
- Studied and implemented algorithms for correspondence estimation and alignment of 3D point clouds.

**MEDIA COVERAGE**

**Reviewed Meta's Text-to-Video Generation AI for MIT Tech Review** Sep 2022
Reported by Melissa Heikkilä from MIT Tech Review

**Article in Seattle Met on AI Models for Art Generation** Sep 2022
Interviewed by Benjamin Cassidy from Seattle Met

**Articles covering GRIT Benchmark** Sep 2022
Our benchmark for evaluation of general purpose vision models featured in several AI news platforms: **MarkTechPost**, **Analytics India Magazine**

**Article on GPV and GRIT in VentureBeat** May 2022
Interviewed by Sharon Goldman from VentureBeat

**GPV on Issue-94 of The Batch** Jun 2021
Work on General Purpose Vision covered in Andrew Ng's popular newsletter - The Batch

**Imagine This! Scripts to Compositions to Videos** Mar 2018
Our text to video generation model featured in several tech media outlets: Engadget, Gizmodo, TNW, Tech Times, Alphr, BGR, Singapore Hardware Zone, Android Headlines

**140K+ YouTube Views on the Demo of our Text to Video Generation Model** Apr 2018
Released as supplementary material for the paper "Imagine This! Scripts to Compositions to Videos"

**TALKS**

**Towards General Purpose Vision Systems** (Oral) Jun 2022
*Venue*: CVPR

**General Robust Image Task Benchmark** (Workshop) Jun 2022
*Venue*: Open World Vision Workshop, CVPR, New Orleans

**General Purpose Vision and the GRIT Benchmark** (Invited) Apr 2022
*Venue*: Microsoft Research, Redmond (Virtual)

**Trends in ML for Vision and Language** (Invited) Feb 2022
*Venue*: IIT Kanpur ML School (Virtual)

**Towards General Purpose Vision Systems** (Invited) Sep 2021
*Venue*: The Computational Linguistics in Québec Consortium (en.cliq-ai.quebec)

**Contrastive Learning for Weakly Supervised Phrase Grounding** (Spotlight) Aug 2020
*Venue*: ECCV 2020

**Representations for Vision and Language** (Invited) Sep 2019
*Venue*: CS598RK: HCI for ML Course, UIUC

**Enhancing Inductive Transfer in Vision-Language Tasks** May 2017
*Venue*: Midwest Vision Workshop, Here Research, Chicago

**Role of Language in Vision** Apr 2017
*Venue*: CS543: Computer Vision Course, UIUC

**Tensorflow Tutorial** (Invited) Mar 2017
*Venue*: Virtusense Technologies

**Deep Learning Panel** (Invited) Nov 2016
*Venue*: Big Data Summit, Research Park, UIUC

**PUBLICATIONS**

**Visual Programming: Compositional visual reasoning without training** arXiv 2022
Tanmay Gupta and Aniruddha Kembhavi

**General Robust Image Task Benchmark** arXiv 2022
Tanmay Gupta, Ryan Marten, Aniruddha Kembhavi, and Derek Hoiem

**Webly Supervised Concept Expansion for General Purpose Vision Models** ECCV 2022
Amita Kamath*, Christopher Clark*, Tanmay Gupta*, Eric Kolve, Derek Hoiem, and Aniruddha Kembhavi

**Towards General Purpose Vision Systems** CVPR 2022
Tanmay Gupta, Amita Kamath, Aniruddha Kembhavi, and Derek Hoiem

**Learning Curves for Analysis of Deep Networks** ICML 2021
Derek Hoiem, Arka Sadhu, Tanmay Gupta, Zhizhong Li, and Michal M. Shlapenthokh-Rothman

**Visual Semantic Role Labeling for Video Understanding** CVPR 2021
Arka Sadhu, Tanmay Gupta, Mark Yatskar, and Aniruddha Kembhavi

**Contrastive Learning for Weakly Supervised Phrase Grounding (Spotlight)** ECCV 2020
Tanmay Gupta, Arash Vahdat, Gal Chechik, Xiaodong Yang, Jan Kautz, and Derek Hoiem

**ViCo: Word Embeddings from Visual Co-occurrences** ICCV 2019
Tanmay Gupta, Alexander Schwing, and Derek Hoiem

**No-Frills Human-Object Interaction Detection: Factorization, Layout Encodings, and Training Techniques** ICCV 2019
Tanmay Gupta, Alexander Schwing, and Derek Hoiem

**Imagine This! Scripts to Compositions to Videos** ECCV 2018
Tanmay Gupta, Dustin Schwenk, Ali Farhadi, Derek Hoiem, and Aniruddha Kembhavi

**Aligned Image-Word Representations Improve Inductive Transfer Across Vision-Language Tasks** ICCV 2017
Tanmay Gupta, Kevin Shih, Saurabh Singh, and Derek Hoiem

**3DFS: Deformable Dense Depth Fusion and Segmentation for Object Reconstruction from a Handheld Camera** arXiv 2016
Tanmay Gupta, Daeyun Shin, Naren Sivagnanadasan, and Derek Hoiem

**Completing 3D Object Shape from One Depth Image** CVPR 2015
Jason Rock, Tanmay Gupta, Justin Thorsen, Junyoung Gwak, Daeyun Shin, and Derek Hoiem

**TECHNICAL BLOG**

**http://bigredt.github.io**
Blog posts about Machine Learning, AI, and Computer Vision.

**AWARDS**

**Outstanding Reviewer** 2021
CVPR 2021

**Outstanding Reviewer** 2020
ECCV 2020

**AI2 Award** 2017
Allen Institute for Artificial Intelligence
Received a $10K award in support of my work on semi-parametric text to video generation

**Sridhar Memorial Prize** 2014
IIT Kanpur
Given to the best student of B.Tech final year Electrical Engineering at the end of 6th Semester

**Todai-IIT Scholarship** (Awarded Twice) 2011 – 2013
University of Tokyo and Mori Seiki Company, Japan
Given to 8 students each from 5 major IITs for academic excellence

**Certificate of Merit for Academic Excellence** (Awarded Thrice) 2010 – 2013
IIT Kanpur, India
Given to top 5% students in each department at IIT Kanpur for academic excellence

**OTHER DUTIES & SERVICES**

**Area Chair for CVPR 2023** Present
Invited to server as an area chair and supervise the entire review process of 20-30 papers

**Reviewer** Jan 2016 – Present
Served as a reviewer for CVPR, ICCV, ECCV, NeurIPS, and TPAMI

**Graduate Teaching Assistant, UIUC** Jan 2017 – May 2017
Prepared and delivered a lecture on Vision-Language research, held discussion hours, and graded homeworks for graduate computer vision course CS543 offered by Prof. Derek Hoiem.

**Organizer, Vision Lunch, UIUC** Aug 2016 – May 2020
Primary computer vision reading and presentation group in the Department of Computer Science, UIUC

**Graduate Research Assistant, UIUC** Aug 2014 – May 2020
*Supervisor*: Prof. Derek Hoiem, Department of Computer Science, UIUC

**Senior Student Research Associate, IIT Kanpur** Aug 2013 – Dec 2013
*Supervisor*: Prof. Aditya K. Jagannatham, Department of Electrical Engineering, IIT Kanpur

**Student Research Associate, IIT Kanpur** Nov 2012 – Apr 2013
*Supervisor*: Prof. Aditya K. Jagannatham, Department of Electrical Engineering, IIT Kanpur

**MENTORING**

**Interns at AI2**
- Ryan Marten (MSc student at UIUC) Summer 2022
- Arka Sadhu (PhD student at USC) Summer 2021
- Arka Sadhu (PhD student at USC) Summer 2020

**Pre-doctoral Young Investigators at AI2**
- Oscar Michel (prev. UG at UChicago) 2022-Present
- Amita Kamath (now a PhD student at UCLA) 2020-2022

**TECHNICAL SKILLS**

**Prog. Languages**: Python, MATLAB, C++, Java
**Deep Learning Frameworks**: Pytorch, Jax / Flax, SeqIO, Pytorch-Lightning, TensorFlow
**Visualization Tools**: Plotly, Tensorboard, Plotly Dash
**Development Tools**: VS Code, Git, SSH, Docker, Emacs, Visual Studio

*Skills and tools used in everyday workflow are underlined.