

# Chris Thomas

## Curriculum Vitae

Virginia Tech  
378 Data and Decision Sciences  
727 Prices Fork Rd  
Blacksburg, VA 24060  
☎ (540) 231-2993  
✉ [chris@cs.vt.edu](mailto:chris@cs.vt.edu)  
🏠 [people.cs.vt.edu/chris](http://people.cs.vt.edu/chris)  
Assistant Professor

### Academic Experience

- 2022- **Assistant Professor**, VIRGINIA TECH, Department of Computer Science, Blacksburg, VA.  
2020-2022 **Postdoctoral Research Scientist**, COLUMBIA UNIVERSITY, Department of Electrical Engineering, New York City, New York. Mentor: Shih-Fu Chang ([link](#)).

### Education

- 2013–2020 **Ph.D. Computer Science**, UNIVERSITY OF PITTSBURGH, Pittsburgh, PA.  
Advisor: Adriana Kovashka. Dissertation Committee: Abhinav Gupta, Adriana Kovashka (Chair), Rebecca Hwa, Diane Litman.  
2009–2013 **B.S. Computer Science**, UNIVERSITY OF PITTSBURGH, Pittsburgh, PA.

### Publications

[Google Scholar \(link\)](#) - 570 citations, h-index 11 as of 08/27/2024

#### Peer-Reviewed Conference Publications

- Notes:** CVPR, ECCV, and ICCV are the top conferences in computer vision with h-5 indices of 422, 238, and 228 respectively. ACL, EMNLP, and NAACL are the top conferences in natural language processing, with h-5 indices of 192, 176, 133. AACL, NeurIPS, ICLR and ICML are top-tier conferences in machine learning generally (h-indices of 212, 309, 303, and 254).
- 2024 ○ Junzhang Liu, Zhecan Wang, Hammad Ayyubi, Haoxuan You, Chris Thomas, Rui Sun, Shih-Fu Chang, Kai-Wei Chang. Detecting Multimodal Situations with Insufficient Context and Abstaining from Baseless Predictions. *To appear*, In Proceedings of the 32nd ACM International Conference on Multimedia (ACM MM). Acceptance rate TBD.
- 2024 ○ Ting-Chih Chen, Chiawei Tang, and Christopher Thomas. MetaSumPerceiver: Multimodal Multi-Document Evidence Summarization for Fact-Checking. Proceedings of the Association for Computational Linguistics (ACL) 2024. Acceptance rate TBD.
- 2024 ○ Alvi Ishmam and Christopher Thomas. Semantic Shield: Defending Vision-Language Models Against Backdooring and Poisoning via Fine-grained Knowledge Alignment. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2024. Seattle, WA. June 2024. Acceptance rate 23.6%.
- 2024 ○ Hammad Ayyubi, Christopher Thomas, Lovish Chum, Rahul Lokesh, Long Chen, Yulei Niu, Xudong Lin, Xuande Feng, Jaywon Koo, Sounak Ray, Shih-Fu Chang. Beyond Grounding: Extracting Fine-Grained Event Hierarchies Across Modalities. The 38th Annual AACL Conference on Artificial Intelligence (AAAI24). Acceptance rate 23.75%.

- 2023 ○ Mingyang Zhou, Yi Fung, Long Chen, Christopher Thomas, Heng Ji, and Shih-Fu Chang. Enhanced Chart Understanding in Vision and Language Tasks via Cross-modal Pre-training on Plot Table Pairs. Findings of the Association for Computational Linguistics: ACL 2023. Acceptance rate 18.4%.
- 2022 ○ Long Chen, Yulei Niu, Brian Chen, Xudong Lin, Guangxing Han, Christopher Thomas, Hammad Ayyubi, Heng Ji, and Shih-Fu Chang. Weakly-Supervised Temporal Article Grounding. Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP), Abu Dhabi. December 2022. Acceptance rate 22.1%.
- 2022 ○ Christopher Thomas, Yipeng Zhang, and Shih-Fu Chang. Fine-Grained Visual Entailment. Proceedings of the European Conference on Computer Vision (ECCV), Tel Aviv, Israel. October 2022. Acceptance rate 28.4%.
- 2021 ○ Brian Chen, Xudong Lin, Christopher Thomas, Manling Li, Shoya Yoshida, Lovish Chum, Heng Ji, and Shih-Fu Chang. Joint Multimedia Event Extraction from Video and Article. Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP2021) Findings. Acceptance rate 25.6%.
- 2021 ○ Yi R. Fung, Christopher Thomas, Revanth Gangi Reddy, Sandeep Polisetty, Heng Ji, Shih-Fu Chang, Kathleen McKeown, Mohit Bansal and Avi Sil. InfoSurgeon: Cross-Media Fine-grained Information Consistency Checking for Fake News Detection. Proceedings of the Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL 2021). **Oral**. Acceptance rate 21.3%.
- 2020 ○ Christopher Thomas and Adriana Kovashka. Preserving Semantic Neighborhoods for Robust Cross-modal Retrieval. Proceedings of the 16th European Conference on Computer Vision (ECCV) 2020. Glasgow, United Kingdom, August 2020. Acceptance Rate: 21.8%.
- 2019 ○ Christopher Thomas and Adriana Kovashka. Predicting the Politics of an Image Using Webly Supervised Data. Proceedings of the Thirty-third Conference on Neural Information Processing Systems (NeurIPS 2019, formerly NIPS). Vancouver, Canada, December 2019. Acceptance Rate: 21.1%.
- 2018 ○ Christopher Thomas and Adriana Kovashka. Artistic Object Recognition by Unsupervised Style Adaptation. Proceedings of the Asian Conference on Computer Vision (ACCV), Perth, Australia, December 2018. Acceptance Rate: 27.9%.
- 2018 ○ Christopher Thomas and Adriana Kovashka. Persuasive Faces: Generating Faces in Advertisements. Proceedings of the British Machine Vision Conference (BMVC), Newcastle upon Tyne, England, September 2018. Acceptance rate 29.9%.
- 2017 ○ Zaeem Hussain, Xiaozhong Zhang, Mingda Zhang, Keren Ye, Christopher Thomas, Zuha Agha, Nathan Ong, and Adriana Kovashka. Automatic Understanding of Image and Video Advertisements. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, Hawaii, July 2017. **Spotlight**. Acceptance rate 8%.
- 2016 ○ Christopher Thomas and Adriana Kovashka. Seeing Behind the Camera: Identifying the Authorship of a Photograph. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, June 2016. Acceptance rate 29.9%.

- 2015 ○ Shi-Kuo Chang, Wen-Hui Chen, Wen-Chyi Lin, and Christopher Thomas. Application of Slow Intelligence Framework for Smart Pet Care System Design. SEKE 2015. Proceedings of the Twenty-Seventh International Conference on Software Engineering and Knowledge Engineering, 74–79.
- 2014 ○ Pamela Jordan, Nancy Green, Christopher Thomas, and Susan Holm. TBI–DOC: Generating Patient & Clinician Reports From Brain Imaging Data. INLG 2014. Proceedings of the Eighth International Natural Language Generation Conference, 12–16.

### Peer-Reviewed Journal Publications

- Notes:** Transactions on Pattern Analysis and Machine Intelligence (T-PAMI) and the International Journal of Computer Vision are top journals in computer vision (h-5 indices of 179 and 88).
- 2022 ○ Desmond U Patton, Nathan Aguilar, Aviv Y Landau, Chris Thomas, Rachel Kagan, Tianai Ren, Eric Stoneberg, Timothy Wang, Daniel Halmos, Anish Saha, Amith Ananthram, and Kathleen McKeown. Community implications for gun violence prevention during co-occurring pandemics: a qualitative and computational analysis study. Preventive Medicine. Pages 107263-107263.
  - 2022 ○ Mesut Erhan Unal, Keren Ye, Mingda Zhang, Christopher Thomas, Adriana Kovashka, Wei Li, Danfeng Qin, and Jesse Berent. Learning to Overcome Noise in Weak Caption Supervision for Object Detection. IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI). June 2022.
  - 2021 ○ Christopher Thomas and Adriana Kovashka. Predicting Visual Political Bias using Webly Supervised Data and an Auxiliary Task. International Journal of Computer Vision (IJCV), 129(11). Pages 2978-3003.
  - 2015 ○ Shi-Kuo Chang, Wen-Hui Chen, Wen-Chyi Lin, and Christopher Thomas. Application of Slow Intelligence Framework for Smart Pet Care System Design. Proceedings of the International Journal of Software Engineering and Knowledge Engineering (IJSEKE 2015), Volume 25, Number 09n10, Pages 1429–1442.

### Peer-Reviewed Workshop Publications

- 2022 ○ Christopher Thomas and Adriana Kovashka. Emphasizing Complementary Samples for Non-Literal Cross-Modal Retrieval. Proceedings of the 5th Multimodal Learning and Applications Workshop, held in conjunction with the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), New Orleans, LA. June 2022. Pages 4632-4641. 8 pages.
- 2016 ○ Christopher Thomas, Adriana Kovashka, Donald Chiarulli and Steven Levitan. A Visual Attention Algorithm Designed for Coupled Oscillator Acceleration. Proceedings of the Twelfth IEEE Embedded Vision Workshop, held in conjunction with the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2016. 8 pages. **Oral.**

### Technical Reports

- 2016 ○ Christopher Thomas. OpenSALICON: An Open Source Implementation of the Salicon Saliency Model. Technical Report. University of Pittsburgh Department of Computer Science.
- 2015 ○ Chris Thomas and Brandon Jennings. Hand Posture’s Effect on Touch Screen Text Input Behaviors: A Touch Area Based Study. arXiv preprint arXiv:1504.02134, 2015.

- 2014 ○ Sean Myers, Timothy Parenti, and Chris Thomas. Student Response Analysis. Technical Report. University of Pittsburgh Department of Computer Science.

### Conference Publications Under Review

- 2024 ○ Chia-Wei Tang, Ting-Chih Chen, Kiet A. Nguyen, Kazi Sajeed Mehrab, Alvi Md Ishmam, Chris Thomas. M3D: MultiModal MultiDocument Fine-Grained Inconsistency Detection. *In submission*, ACL Rolling Review, June 2024.
- 2024 ○ Zhecan Wang, Junzhang Liu, Chia-Wei Tang, Hani Alomari, Anushka Sivakumar, Rui Sun, Wenhao Li, Md. Atabuzzaman, Hammad Ayyubi, Haoxuan You, Alvi Md Ishmam, Kai-Wei Chang, Shih-Fu Chang, Chris Thomas. JourneyBench: A Challenging One-Stop Vision-Language Understanding Benchmark of Generated Images. *In submission*.

---

### Awards

- 2023 CVPR 2023 Outstanding Reviewer (top 5% of experienced reviewers)
- 2021 ICCV 2021 Outstanding Reviewer (top 5% of experienced reviewers)
- 2021 CVPR 2021 Outstanding Reviewer
- 2020 ECCV 2020 Outstanding Reviewer
- 2019 NeurIPS 2019 Travel Grant
- 2019 ICCV 2019 Outstanding Reviewer
- 2019 CVPR 2019 Outstanding Reviewer
- 2018 CS50 Merit Pre-Doctoral Fellowship (Andrew Mellon Fellowship)
- 2016 Orrin E. and Margaret M. Taulbee Award (recognizes teaching excellence and academic progress)
- 2016 CS Department Best Research Poster Runner-Up. [Click here to download.](#)
- 2015 CS Department Teaching Assistant Award
- 2015 CS Department Best Research Poster Award. [Click here to download.](#)
- 2014 CS Department Teaching Assistant Award
- 2014 CS Department Best Research Poster Award. [Click here to download.](#)
- 2014 University of Pittsburgh Elizabeth Baranger Excellence In Teaching Award
- 2014 Kenneth P. Dietrich School of Arts and Sciences Postgraduate Fellowship

---

### Selected Press

- 2023 “When Trust is a Factor” Grant Currin, Columbia Engineering Magazine, “The Information Issue.” Fall 2023, vol. 65 no. 1. [Click here to visit.](#)
- 2021 “Who should stop unethical A.I.?” Matthew Hutson, February 15, 2021. The New Yorker. [Click here to visit.](#)
- 2018 “Using machine learning to generate persuasive faces for ads.” August 6, 2018. [Click here to visit.](#)
- 2015 “Who’s behind the camera: computer vision and author identification.” August 24, 2015. [Click here to visit.](#)
- 2015 “New photographs from dead photographers with convolutional neural networks.” November 12, 2015. [Click here to visit.](#)

---

## Professional Service

### Area Chair / Meta-reviewer

WACV (*IEEE Winter Conference on Applications of Computer Vision*) - 2024

### Reviewing

**Notes:** All conference reviewing below indicates Program Committee membership.

ACL ARR (*Association for Computational Linguistics Rolling Review*) - 2023, 2024

AAAI (*Association for the Advancement of Artificial Intelligence Conference on Artificial Intelligence*) - 2020, 2021, 2022, 2023, 2024

BMVC (*British Machine Vision Conference*) - 2019

CVPR (*IEEE Conference on Computer Vision and Pattern Recognition*) - 2019 (Outstanding reviewer), 2020, 2021 (Outstanding reviewer), 2022, 2023 (Outstanding reviewer), 2024

DIRA (*Diagram Image Retrieval and Analysis Workshop held at CVPR2020*) - 2020

ECCV (*European Conference on Computer Vision*) - 2020 (Outstanding reviewer), 2024

ICCV (*IEEE International Conference on Computer Vision*) - 2019 (Outstanding reviewer), 2021 (Outstanding reviewer), 2023

ICLR (*International Conference on Learning Representations*) - 2022, 2023, 2024

ICMI (*ACM International Conference on Multimodal Interaction*) - 2019

ICML (*International Conference on Machine Learning*) - 2024

IJCV (*International Journal of Computer Vision*) - 2020

ICVGIP (*Indian Conference on Computer Vision, Graphics and Image Processing*) - 2016

LUV (*Learning from Unlabeled Videos (CVPR Workshop)*) - 2019

NeurIPS (*Conference on Neural Information Processing Systems*) - 2020, 2021, 2023, 2024

T-PAMI (*IEEE Transactions on Pattern Analysis and Machine Intelligence*) - 2022

WACV (*IEEE Winter Conference on Applications of Computer Vision*) - 2020

### Departmental Service

Graduate Admissions Committee - 2023-

### Synergistic Activities

CS Department Research Poster Judge - 2024

TechGirls - 2024-

BEE VT - 2024-

---

## Teaching

**Notes:** All SPOT survey response scores indicate the average of all questions related to instructor.

### Graduate Courses Taught

- o CS 5814: Introduction to Deep Learning
  - Semesters taught: Spring 2024
  - SPOT Evaluation: TBD

- CS 5864: Learning-based Computer Vision - Developed and taught new course
  - Semesters taught: Fall 2023 (enrollment: 28)
  - SPOT Evaluation: 5.29 (2023)
- CS 6804: Multimodal Vision
  - Semesters taught: Spring 2023 (enrollment: 18)
  - SPOT Evaluation: 5.7 (2023)
- CS 5974: Independent Study
  - Semesters taught: Spring 2023 (enrollment: 1)

### Undergraduate Courses Taught

- CS 4994: Undergraduate Research
  - Semesters taught: Fall 2023 (enrollment: 1, for capstone credit)
- CS 2974: Independent Study
  - Semesters taught: Spring 2024 (enrollment: 1)

## Students Supervised

### PhD Students Supervised

- 2024- Md. Atabuzzaman
- 2023- Hani Alomari
- 2022- Eslam Hussein
- 2022- Alvi Ishmam

### Masters Students Supervised

- 2024- Zaber Hakim
- 2024- Najibul Sarker
- 2024- Aafiya Hussain
- 2024- Anushka Sivakumar
- 2023-2024 Hanwen Zhang (Graduated)
- 2023- Andrew Zhang
- 2022-2023 Ting-Chih Chen (Graduated)
- 2022-2024 Chiawei Tang (Graduated)

### Supervised as Postdoc

- 2021-2022 Thomas Zollo (Masters)
- 2021-2022 Yipeng Zhang (Masters)
- 2021 Xinyu Wei (Undergraduate)
- 2020 Adam Svystun (Masters)
- 2020-2021 Catherine Chen (Undergraduate)
- 2020-2022 Timothy Wang (Undergraduate)
- 2020-2022 Daniel Halmos (Undergraduate)

---

## Committee Memberships

**Notes:** *This list is currently incomplete.*

### PhD Student Committee Membership

Jingyuan Qi  
Haisam Rafid  
Eman Abdelrahman  
Michael Goldsworthy  
Barry Yao  
Xiaona Zhou  
Afrina Tabassum  
Kiet Nguyen  
Amun Kharel  
Chase Vickery  
Anish Narkar  
Medha Sawhney  
Hidir Yesiltepe  
Kiymet Akdemir  
Morteza Haji  
Mingsi Liao

### Master's Student Committee Membership

Brendan Arnold (Biomedical Engineering)  
Daniel Stover (Electrical and Computer Engineering)  
Sachiko Sakai (Geosciences)  
Vedant Shah  
Swapnil Singh  
Amartya Dutta  
Tuna Meral  
Yusuf Dalva

---

## Invited Talks

- 2024 Google Research. Perception Group (DeepMind). April 24, 2024.
- 2023 University of Missouri. Colloquium (Host: Henry Kang). Dec 1, 2023.
- 2022 Iowa State University. 2022 Midwest Big Data Summer School (Host: Christopher Quinn). May 19, 2022.
- 2022 Florida State University. Colloquium (Host: Xiuwen Liu). February 11, 2022.
- 2021 Binghamton University. Departmental Colloquium (Host: Weiyi Meng), December 13, 2021.
- 2021 Iowa State University. Departmental Colloquium (Host: Yan-Bin Jia), December 3, 2021.
- 2021 Columbia University, School of Engineering and Applied Sciences' Colloquium for Engineering Students (Host: David Vallancourt), November 12, 2021.

## Selected Projects and Code

- Github **Visual Entailment** – Project page containing code and other materials for our *Fine-Grained Visual Entailment* ECCV 2022 paper.
- Project Page **Semantic Neighborhoods** – Project page containing code and other materials for our *Preserving Semantic Neighborhoods for Robust Cross-modal Retrieval* ECCV 2020 paper.
- Project Page **Politics** – Project page containing dataset, code, and other materials for our *Predicting the Politics of an Image Using Webly Supervised Data* NeurIPS 2019 paper.
- Project Page **Artistic Objects** – Project page containing dataset and other materials for our *Artistic Object Recognition by Unsupervised Style Adaptation* ACCV 2018 paper.
- Project Page **Photographer Identification** – Project page containing dataset, downloadable models, and more for our *Seeing Behind the Camera: Identifying the Authorship of a Photograph* CVPR 2016 paper.
- GitHub **OpenSALICON** – An open source implementation of the SALICON saliency algorithm containing trained saliency models and a modified version of the Caffe deep learning library.
- Download **My DBMS** – A toy implementation of SAP HANA's L1-delta and L2-delta from *this* paper.

## Languages

- English **Mothertongue**
- Spanish **Moderate**

*Basic conversational ability*