

Jack Hessel

Email: jmhessel@gmail.com

Website: www.jmhessel.com

Twitter: [@jmhessel](https://twitter.com/jmhessel)

Education

Ph.D., Computer Science

Cornell University

August 2014 - August 2020

Advised by Professor Lillian Lee

Natural Language Processing and Machine Learning

Thesis: Learning from Multimodal Web Data

B.A., Computer Science and Math/Statistics

Carleton College

September 2010 - June 2014

Magna Cum Laude

Honors in Computer Science

Publications

Refereed Publications

Hyunwoo Kim, **Jack Hessel**, Liwei Jiang, Ximing Lu, Youngjae Yu, Pei Zhou, Ronan Le Bras, Malihe Alikhani, Gunhee Kim, Maarten Sap, Yejin Choi. "SODA: Million-scale Dialogue Distillation with Social Commonsense Contextualization." EMNLP 2023.

Amita Kamath, **Jack Hessel**, Kai-Wei Chang. "Text encoders are performance bottlenecks in contrastive vision-language models." EMNLP 2023.

Amita Kamath, **Jack Hessel**, Kai-Wei Chang. "What's "up" with vision-language models? Investigating their struggle to understand spatial relations." EMNLP 2023.

Seungju Han, Junhyeok Kim, **Jack Hessel**, Liwei Jiang, Jiwan Chung, Yejin Son, Yejin Choi, Youngjae Yu. "Reading Books is Great, But Not if You Are Driving! Visually Grounded Reasoning about Defeasible Commonsense Norms." EMNLP 2023.

Peter West, Ronan Le Bras, Taylor Sorensen, Bill Yuchen Lin, Liwei Jiang, Ximing Lu, Khyathi Chandu, **Jack Hessel**, Ashutosh Baheti, Chandra Bhagavatula, Yejin Choi. "NovaCOMET: Open Commonsense Foundation Models with Symbolic Knowledge Distillation." Findings of EMNLP 2023.

Jae Sung Park, **Jack Hessel**, Khyathi Chandu, Paul Pu Liang, Ximing Lu, Qiuyuan Huang, Peter West, Jianfeng Gao, Ali Farhadi, Yejin Choi. "Localized Symbolic Knowledge Distillation for Visual Commonsense Models." NeurIPS 2023.

Wanrong Zhu*, **Jack Hessel***, Anas Awadalla, Samir Yitzhak Gadre, Jesse Dodge, Alex Fang, Youngjae Yu, Ludwig Schmidt, William Yang Wang, Yejin Choi. "Multimodal C4: An Open, Billion-scale Corpus of Images Interleaved With Text." NeurIPS Datasets+Benchmarks 2023.

Yonatan Bitton*, Hritik Bansal*, **Jack Hessel***, Rulin Shao, Wanrong Zhu, Anas Awadalla, Josh Gardner, Rohan Taori, Ludwig Schimdt. "VisIT-Bench: A Benchmark for Vision-Language Instruction Following Inspired by Real-World Use." NeurIPS Datasets+Benchmarks 2023.

Yizhong Wang, Hamish Ivison, Pradeep Dasigi, **Jack Hessel**, Tushar Khot, Khyathi Raghavi Chandu, David Wadden, Kelsey MacMillan, Noah A. Smith, Iz Beltagy, Hannaneh Hajishirzi. "How Far Can Camels Go? Exploring the State of Instruction Tuning on Open Resources." NeurIPS Datasets+Benchmarks 2023.

Seungju Han, **Jack Hessel**, Nouha Dziri, Yejin Choi, Youngjae Yu. "CHAMPAGNE: Learning Real-world Conversation from Large-Scale Web Videos." ICCV 2023.

Nitzan Bitton-Guetta*, Yonatan Bitton*, **Jack Hessel**, Ludwig Schmidt, Yuval Elovici, Gabriel Stanovsky, Roy Schwartz. "Breaking Common Sense: WHOOPS! A Vision-and-Language Benchmark of Synthetic and Compositional Images." ICCV 2023.

Jack Hessel, Ana Marasović, Jena D. Hwang, Lillian Lee, Jeff Da, Rowan Zellers, Robert Mankoff, Yejin Choi. "Do Androids Laugh at Electric Sheep? Humor 'Understanding' Benchmarks from The New Yorker Caption Contest." ACL 2023; *Best Paper Award, one of 3 / 3872 submissions*

Liunian Harold Li, **Jack Hessel**, Youngjae Yu, Xiang Ren, Kai-Wei Chang and Yejin Choi. "Symbolic Chain-of-Thought Distillation: Small Models Can Also 'Think' Step-by-Step." ACL 2023.

Youngjae Yu, Jiwan Chung, Heeseung Yun, **Jack Hessel**, JaeSung Park, Ximing Lu, Prithviraj Ammanabrolu, Rowan Zellers, Ronan Le Bras, Gunhee Kim, Yejin Choi. "Multimodal Knowledge Alignment with Reinforcement Learning." CVPR 2023.

Rajkumar Ramamurthy, Prithviraj Ammanabrolu, Kianté Brantley, **Jack Hessel**, Rafet Sifa, Christian Bauckhage, Hannaneh Hajishirzi, and Yejin Choi. "Is Reinforcement Learning (Not) for Natural Language Processing?: Benchmarks, Baselines, and Building Blocks for Natural Language Policy Optimization." ICLR 2023 (spotlight).

Ximing Lu, Sean Welleck, Liwei Jiang, **Jack Hessel**, Lianhui Qin, Peter West, Prithviraj Ammanabrolu, and Yejin Choi. "Quark: Controllable Text Generation with Reinforced Unlearning." NeurIPS 2022 (Oral; top 1% of submissions).

Jack Hessel*, Jena D. Hwang*, Jae Sung Park, Rowan Zellers, Chandra Bhagavatula, Anna Rohrbach, Kate Saenko, and Yejin Choi. "The Abduction of Sherlock Holmes: A Dataset for Visual Abductive Reasoning." ECCV 2022 (oral).

Sarah Wiegrefe, **Jack Hessel**, Swabha Swayamdipta, Mark Riedl, and Yejin Choi. "Reframing Human-AI Collaboration for Generating Free-Text Explanations." NAACL 2022.

Yanpeng Zhao, **Jack Hessel**, Youngjae Yu, Ximing Lu, Rowan Zellers, and Yejin Choi. "Connecting the Dots between Audio and Text without Parallel Data through Visual Knowledge Transfer." NAACL 2022.

Peter West, Chandra Bhagavatula, **Jack Hessel**, Jena D. Hwang, Liwei Jiang, Ronan Le Bras, Ximing Lu, Sean Welleck, and Yejin Choi. "Symbolic Knowledge Distillation: from General Language Models to Commonsense Models." NAACL 2022.

Rowan Zellers, Jiasen Lu, Ximing Lu, Youngjae Yu, Yanpeng Zhao, Mohammadreza Salehi, Aditya Kusupati, **Jack Hessel**, Ali Farhadi, and Yejin Choi. "MERLOT Reserve: Neural Script Knowledge through Sound, Language, and Vision." CVPR 2022.

Rowan Zellers*, Ximing Lu*, **Jack Hessel***, Youngjae Yu, Jae Sung Park, Jize Cao, Ali Farhadi, and Yejin Choi. "MERLOT: Multimodal Neural Script Knowledge Models." NeurIPS 2021 (Oral; top 1% of submissions).

Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. "CLIPScore: A Reference-free Evaluation Metric for Image Captioning." EMNLP 2021.

Jack Hessel and Alexandra Schofield. "How effective is BERT without word ordering? Implications for language understanding and data privacy." ACL 2021.

Jack Hessel and Lillian Lee. "Does my multimodal model learn cross-modal interactions? It's harder to tell than you might think!" EMNLP 2020.

Jack Hessel, Zhenhai Zhu, Bo Pang, and Radu Soricut. "Beyond Instructional Videos: Probing for More Diverse Visual-Textual Grounding on YouTube." EMNLP 2020.

Gregory Yauney, **Jack Hessel**, and David Mimno. "Domain-Specific Lexical Grounding in Noisy Visual-Textual Documents." EMNLP 2020.

Jack Hessel, Lillian Lee, and David Mimno. "Unsupervised Discovery of Multimodal Links in Multi-image, Multi-sentence Documents." EMNLP 2019.

Jack Hessel, Bo Pang, Zhenhai Zhu, and Radu Soricut. "A Case Study on Combining ASR and Visual Features for Generating Instructional Video Captions" CoNLL 2019.

Jack Hessel and Lillian Lee. "Something's Brewing! Early Prediction of Controversy-causing Posts from Discussion Features." NAACL 2019.

Jack Hessel, David Mimno, and Lillian Lee. "Quantifying the Visual Concreteness of Words and Topics in Multimodal Datasets." NAACL 2018.

Jack Hessel, Lillian Lee, and David Mimno. "Cats and Captions vs. Creators and the Clock: Comparing Multimodal Content to Context in Predicting Relative Popularity" WWW 2017.

Jack Hessel, Chenhao Tan and Lillian Lee. "Science, AskScience and BadScience: On the Coexistence of Highly Related Communities." ICWSM 2016.

Jack Hessel, and Sherri Goings. "Using Reproductive Altruism to Evolve Multicellularity in Digital Organisms." ECAL 2013.

Refereed Workshop Publications

Jack Hessel, and David Mimno. "Aligning Images and Text in a Digital Library." Computer Vision in Digital Humanities Workshop @ DH 2017.

Jack Hessel, Alexandra Schofield, Lillian Lee, and David Mimno. "What do Vegans do in their Spare Time? Latent Interest Detection in Multi-Community Networks." Networks Workshop @ NeurIPS 2015.

Jack Hessel, Nicolas Savva, and Kimberly J Wilber. "Image Representations and New Domains in Neural Image Captioning." Vision/Language Workshop @ EMNLP 2015.

Albright, Evan, **Jack Hessel**, Nao Hiranuma, Cody Wang, and Sherri Goings. "A Comparative Analysis of Popular Phylogenetic Reconstruction Algorithms." MICS 2014.

Employment

Research

Research @ Samaya AI Oct 2023 – present

Research Scientist, *Allen Institute for Artificial Intelligence* July 2021 – Oct 2023

Postdoctoral Young Investigator, *Allen Institute for Artificial Intelligence* Oct 2020 – July 2021

Host: Yejin Choi.

Working on commonsense reasoning, vision and language, etc.

Research Intern, *Google Research* Summer 2019, Summer 2018

Hosts: Bo Pang and Zhenhai Zhu.

Worked with the natural language understanding team on video language joint learning: this work was published at *CoNLL, 2019* and *EMNLP 2020*.

Research Intern, *Facebook, Inc.* Summer 2017

Host: Amit Bahl.

Worked with the Core Data Science team on personalized language modeling, and cross-modal retrieval.

Research Intern, Twitter, Inc. Summer 2016
 Host: Clément Farabet.
 Worked with the Cortex Team as their first intern on large-scale/multimodal node embeddings in graphs, language modeling, and engagement prediction.

Research Intern, Washington University, St. Louis REU Summer 2013
 Host: Kilian Weinberger.
 Contributed to a GPU support vector machine package that accompanies Tyree et al.'s "Parallel Support Vector Machines in Practice." 2014.

Teaching

Invited Visiting Instructor, Computer Science Dept., Carleton College Spring 2019
 Lead instructor of Natural Language Processing and Mathematics of Computer Science; 30+ students in each class

Teaching Assistant, Cornell University Various
 Language and Information, 2016; Machine Learning for Data Science, 2015; Intro to Computer Graphics, 2014

Invited Talks, Academic Service, Honors, etc.

Invited Talks

Carnegie Mellon University: "Hurdles for vision+language models." 2023

Cornell University: "Hurdles for vision+language models." 2023

NICE CVPR Workshop: "Recognition isn't enough for image captioning." 2023

Guest lecture for Caltech CS148 w/ Yejin Choi "Common Sense: The Dark Matter of Language Intelligence" 2023

Guest lecture for Caltech CS148 w/ Yejin Choi "Common Sense: The Dark Matter of Language Intelligence" 2023

The AI Talks (NUS/NTU Singapore): "The Case for Reasoning Beyond Recognition" 2023

Procter & Gamble: "The Case for Reasoning Beyond Recognition" 2022

Seoul National University: "The Case for Reasoning Beyond Recognition" 2022

Adobe Research: "The Case for Reasoning Beyond Recognition" 2022

University of Washington: Two Lectures (Linguistics, Honors Col.) "New Frontiers in Multimodal Grounding" 2022

University of Pittsburgh: "(at least) Two Conceptions of Visual-Textual Grounding" 2020

Allen Institute for AI: "The Promise and Perils of Learning Grounding from Visual-Textual Web Data." 2020

UNC Chapel Hill: "The Promise and Perils of Learning Grounding from Visual-Textual Web Data." 2020

Rutgers University: "Multimodal Grounding from User-generated Web Content." 2019

SRI International: "Multimodal Grounding from User-generated Web Content." 2019

Cornell University: PhD Colloquium, "Unsupervised Learning From Multimodal Documents." 2019

University of Pittsburgh: "Grounding Images from a Digital Library in their Textual Contexts." 2018

Cornell University: Two Guest Lectures for CS4300, "Practical Unsupervised Learning" 2015

Carleton College: "The Role of Altruism on Kickstarter" 2014

Program Committees/Reviewing/Area Chairing

NB: As of 2023, I don't keep this section up-to-date; I review/area chair for many major NLP/CV/ML conferences, and occasionally journals as well.

Conference/Journal Review Committees

ACL 2016, 2017, 2018, 2019, 2020, 2021	ICML 2020, 2021
EMNLP 2017, 2018, 2019, 2020, 2021, 2022	NeurIPS 2021, 2022
NAACL 2018, 2019, 2021	ICLR 2021, 2022
ACL Rolling review: reviewing/action editing for most cycles since inception	AAACL 2020
	CoNLL 2019, 2020, 2021, 2022 (publicity chair)
	JAIR 2020

PLOS One 2020
ICWSM 2018
EACL 2017, 2021
AAAI 2017, 2022 (Area chair)
EAAMO 2021, 2022 (Area chair)

Student Research Workshop @ EACL: 2021
Noisy User-generated Text @ EMNLP, COLING: 2018,
2019, 2020, 2021, 2022
Practical ML 4 the Developing World @ ICLR: 2020
BlackAIR 2021

Workshop/Grant Reviewer

Black in AI @ NeurIPS: 2017, 2018, 2019, 2020
Student Research Workshop @ NAACL: 2018
Student Research Workshop @ ACL: 2020, 2021

Honors

ACL Best Paper Award	2023
Top Reviewer Recognition	Various
ACL 2018+2020+2021, NAACL 2019, EMNLP 2018+2019, CoNLL 2019, ICML 2020	
Pitt Digital Humanities graduate speaker series speaker	2018
MICS Conference Best Paper Award	2014
Phi Beta Kappa, Beta of Minnesota	2014
Sigma Xi Inductee	2014

Volunteer Service

Lesson Planner and Volunteer, *Expand Your Horizons @ Cornell* Spring 2015, 2014
Planned for and taught at a one-day conference for 7th-9th grade girls to encourage interest in math and science.

Volunteer Elementary School Teacher, Fall 2014
Volunteered once per week teaching 2-5th grade students how to code using code.org

Development Experience

Open Source Contributions

Developed a TreeLSTM in TensorFlow2: this neural network dynamically changes its topology on a per-example basis (https://github.com/jmhessel/recursive_nn_tf2)

Developed fmpytorch in 2017 (150+ stars on GitHub): a cythonized implementation of second order factorization machines in pytorch (<https://github.com/jmhessel/fmpytorch>)

Developed fightingwords in 2015 for comparing word usage rate differences between corpora; used in several refereed publications (<https://github.com/jmhessel/FightingWords>)

Pull requests merged to *Keras*, and *Gensim*, and *tensorflow*

Technical Skills

Machine Learning Skills: Various machine learning/statistical toolkits/languages (e.g. sklearn, Tensorflow, PyTorch, R, etc.). Experience working with large, multi-faceted datasets.

Development Skills: Object-oriented programming (Python, Java, C++), parallel programming experience on CPUs + GPUs + TPUs, experience with various languages, development environments, version control systems, operating systems.

References

Available upon request