

# Eric Chu

[echu@mit.edu](mailto:echu@mit.edu)  
<https://web.media.mit.edu/~echu>

## Education

- 2017–2021 **PhD, Massachusetts Institute of Technology (MIT)**  
Advisor: Deb Roy, Jacob Andreas  
Thesis: Learning human beliefs with language models
- 2015–2017 **MS, Massachusetts Institute of Technology (MIT)**  
Thesis: Audio-visual sentiment analysis for learning emotional arcs in movies  
Committee: Deb Roy, Hugo Larochelle, Pattie Maes
- 2010–2014 **BS, University of California, Berkeley**  
Electrical Engineering & Computer Science
- 2013 Visiting Student, **University of Oxford**  
Mathematics and Bioinformatics

## Research Interests

I work primarily on making machine learning and natural language processing models more *capable*, *flexible*, and *deployable*, with the ultimate goal to study and help guide human health and behavior. I have experience and interest in generative sequential models across different modalities, external memory networks for better adaptability, social reasoning in NLP, and robustness and human-AI interaction. My work sometimes intersects with computational social science; there I am particularly motivated by applications in the way humans communicate – such as the use of stories, pragmatics, and mass media.

## Industry Research Experience

- 2019 **Facebook AI Research (FAIR)**, *Research Intern*  
Hosts: Jason Weston, Stephen Roller
- Worked on dialogue generation, experimenting with large-scale (1.5B parameter) memory models that blend retrieval and generation.
- 2018 **Google Brain**, *Research Intern*  
Host: Peter J. Liu
- Developed the first end-to-end model for unsupervised, multi-document, abstractive summarization. Paper accepted to ICML.

## Academic Research Experience

2015–2021 **MIT**, *PhD Candidate at the Media Lab*

Projects in tying semantic modeling to human outcomes across machine learning, natural language processing, computer vision, audio and speech processing, computational social science, and human-computer interaction.

- Novel approach for probing language models in order to predict media effects; validated in COVID-19 and other settings. Preparing for submission to general-purpose journal.
- Deep generative sequential models: the first end-to-end unsupervised, multi-document, abstractive summarization model accepted to ICML; natural-language plans to guide trajectories, style transfer in text, text simplification.
- External memory networks for domain adaptation and learning personas: new architecture allows incorporation of outside knowledge, and is useful for applications in hate speech detection and narrative analysis. Work accepted to EMNLP and ACL.
- Master's thesis: novel neural audio-visual sentiment models, used to create emotional arcs in movies, which are then shown to be significant predictors of engagement.

2011–2014 **UC Berkeley**, *Undergraduate researcher*

Worked on projects across data analysis, computer vision, and biomedical imaging.

- “Modeling topological sculptures”: created GUI interface to construct 3D-printable surfaces. Researched 3D reconstruction techniques from photos with limited and obstructed view.
- “Topological data analysis”: initiated group study course on TDA, an application of algebraic topology that analyzes different dimensional holes in data.
- “Effect of varying bio-particles on magnetic particle imaging resolution”: wrote Matlab code and ran wet lab experiments in the Berkeley Imaging Systems Lab.

2013 **Oxford University**, *Undergraduate researcher*

- “Support vector machines to classify protein behavior from A-domain sequences” in Jotun Hein's Evolutionary Bioinformatics group.

---

## Additional Work Experience

2015 **Facebook**, *Data Scientist* Ads Integrity Team

Worked on ad moderation, with a focus on the human-AI system

- Led multi-office (Menlo Park, Austin, Hyderabad) project running A/B tests on how to present model predictions and uncertainty to human reviewers, leading to gains in precision and recall.
- Prototyped new document-based models, incorporating structural and semantic features.

2014 **Facebook**, *Data Scientist Intern*, Groups Team

Analyzed community and interest groups, focused on health and emergent behavior

- Built a classifier for detecting organic marketplace groups, and helped inform the creation of Facebook Marketplace in Feb 2015.
- Helped set team direction to focus on smaller-sized groups in international markets.

2013 **Knewton**, *Software Engineer Intern*

- Full-stack development of data storage, file transfer pipeline, error handling, and dashboard for education-tech startup.

---

## Teaching Experience

- 2020 **MAS.S62: Understanding Public Thought** , TA  
Designed labs to introduce applied machine learning and NLP for public opinion mining.
- 2018 **MAS.S10 AI & Equality**, *Organizer, TA*  
Co-organized and led discussions in MIT seminar course on ethical implications of AI – specifically how AI can both promote and impede equality.

---

## Selected Publications

- [1] **Eric Chu**, Jacob Andreas, Steve Ansolabehere, Deb Roy. Neural language models of media consumption can predict public opinion. *In submission to Nature Human Behavior*.
- [2] Nabeel Gillani, **Eric Chu**, Doug Beeferman, Rebecca Eynon, and Deb Roy. Parents' Online School Reviews Reflect Several Racial and Socioeconomic Disparities in K–12 Education. *AERA Open*, 7, 2332858421992344. 2021.
- [3] **Eric Chu**, Deb Roy, and Jacob Andreas. "Are visual explanations useful? a case study in model-in-the-loop prediction." *In submission*.
- [4] Prashanth Vijayaraghavan, **Eric Chu**, and Deb Roy. "DAPPER: Learning Domain-Adapted Persona Representation Using Pretrained BERT and External Memory." Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing. 2020.
- [5] **Eric Chu**, Nabeel Gillani, and Sneha Priscilla Makini. "Games for fairness and interpretability." Companion Proceedings of the Web Conference. 2020. (Also accepted to ICLR'20, Towards Trustworthy ML Workshop).
- [6] **Eric Chu**\* and Peter Liu\*. "Meansum: a neural model for unsupervised multi-document abstractive summarization." International Conference on Machine Learning. PMLR, 2019.
- [7] **Eric Chu**\*, Prashanth Vijayaraghavan\*, and Deb Roy. "Learning Personas from Dialogue with Attentive Memory Networks." Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing. 2018.
- [8] **Eric Chu** "Artistic Influence GAN." AI for Creativity and Design, Neurips Workshop. 2018.
- [9] **Eric Chu** and Deb Roy. "Audio-visual sentiment analysis for learning emotional arcs in movies." 2017 IEEE International Conference on Data Mining (ICDM). IEEE, 2017.

---

## Service

Conference Reviewer: ICML '19, ACL '20, NeuRIPS '20

---

## Grants

Samsung, \$150K: Automatic Feature Learning for Video Content

## Selected Press Coverage

- MIT* Study shows online school reviews reflect school demographics more than effectiveness [\[link\]](#)
- Variety* A Team of MIT Scientists Taught an AI to Get Emotional Over Movies [\[link\]](#)
- Tech Radar* How artificial intelligence is creating new ways of storytelling [\[link\]](#)
- NYTimes* The Remote Control, Out of Control: Why à la Carte TV Is Too Much for a Trekkie [\[link\]](#)