

## Zachary Bamberger

Azulai 3, Apt. 12,  
Ramat Alon, Haifa 3298903  
(054) 256-6993  
zacharybamberger1@gmail.com

## Skills

Natural Language Processing, Machine Learning, Artificial Intelligence, Reinforcement Learning, Multi-Armed Bandits, Machine Learning for Human Behaviour, Algorithms, Object Oriented Programming, Distributed Computing, Python, C++, Tensorflow, Numpy, Pandas, Scikit-Learn, Keras, Google Cloud Platform (GCP),

## Experience

September 2020 - October 2021

### **Google, Mountain View, CA** - *Software Engineer*

- Built a distributed computing pipeline to introduce realistic synthetic typos to existing datasets. I used this pipeline to create datasets which test language model robustness to typo noise.
- Implemented a distributed computing pipeline to unbalance federated datasets, thereby enabling more realistic simulations of federated language models.
- Created a dashboard to display metrics from live experiments with language models. This dashboard refreshes in fixed sub-daily intervals and saves a team of 40+ engineers from running daily operations that take at least 10 minutes to run.

July 2019 - September 2020

### **Google, Mountain View, CA** - *Engineering Resident*

- Implemented an on-device transformer for language modeling tasks (e.g., next word and emoji predictions for GBoard) that is 10x smaller than BERT-Small.
- Developed an offline reinforcement learning library for Google News and YouTube recommender systems. I implemented offline DQN and Policy Gradient (REINFORCE) estimators (tested via Markov Decision Processes).

January 2018 - August 2018

### **Intel, Haifa, Israel** - *Software Engineering Student*

- Optimized common operations (e.g. convolutions) in deep learning frameworks (via x86 assembly and C++) for both current and future Intel processors.
- Worked towards decreasing cycle count and power consumption for various deep learning topologies on Intel processors (e.g. IRV2, SRGAN, and NMT).
- Wrote Python scripts to parse and analyze the projected performance of various neural networks.

## Education

October 2021 - Current

**Technion, Haifa, Israel** – *Computer Science MSc (with thesis)*

*Notable Courses:* Seminar in Natural Language Processing, Advanced topics in Deep Learning, Behavioral Machine Learning

*Awards:* Graduate scholarship covering full tuition plus stipend.

*Teaching:* Introduction to Machine Learning.

*Research:*

- Introduce a novel collection of probing objectives to measure the extent to which model embeddings are able to capture argumentative linguistic properties.
- Design architectural components to make pre-trained language models more effective at evaluating argument persuasiveness as well as relevant discourse-oriented probing tasks.
- Model persuasiveness with graph neural networks (GNNs) that process knowledge graph representations of arguments.

*Open Source Projects:*

- Semi-Supervised Learning for Computer Vision ([GitHub](#)).
- NLP Seminar Report on Persuasive Argumentation ([PDF](#))

August 2015 - May 2019

**Cornell University, Ithaca, New York** – *Computer Science BSc*

*GPA:* 3.315 (Cumulative), 3.453(Major)

*Notable Courses:* Graduate Natural Language Processing, Graduate Artificial Intelligence, Graduate Analysis of Algorithms, Machine Learning for Intelligent Systems, Natural Language Processing

*Awards:* College of Engineering's Dean's List for Spring 2017 Semester (achieved GPA > 3.5)

*Teaching:* Natural Language Processing, Operating Systems

*Research:*

- Constructed a new dataset for argument mining and persuasion from the existing UKP dataset.
- Predicted argument persuasiveness score from text using BERT and analysis of structural features of arguments via Random Forest.
- Utilized an argument mining tool from Argument Mining with Structured SVMs and RNNs (Niculae, Park, and Cardie 2017) to construct and analyze knowledge graph representations of arguments.

## Volunteering

NLP Research Night (March 2022)

- Organized the first NLP Research Night event at the Technion, which hosted ~60 undergrads, ~60 graduate students, 10 faculty members, and 40 industry researchers.
- Raised ~20,000 NIS for the event from Grove Ventures, OurCrowd, and Technion MLIS.
- We discussed some of the most pressing NLP research questions in a round table setting. Each participating company (Facebook, Google, IBM, Amazon, Gong, and Intel) moderated discussions at their tables. Through this event we brought together some of Israel's brightest NLP researchers in an open and casual setting.

Student Council Representative (October 2021 - Current)

- Elected as the Graduate Student representative for the Computer Science faculty at the Technion.
- Organized and volunteered in both academic and social events for Computer Science students.

TreeHacks Mentor (February, 2020)

- Machine learning mentor for Stanford's annual hackathon in.