

Belén Carolina Saldías Fuentes

Cambridge, MA, USA | belen@mit.edu
belencarolina.com | github.com/bcsaldias

(as of June 2024)

EDUCATION

Massachusetts Institute of Technology (MIT) – Cambridge, MA, USA

Ph.D. Candidate.

September 2018 – present

- Research advisor: Professor Deb Roy.
- Topic: Leveraging Large Language Models (LLMs) to reimagine opportunities for purposeful online spaces, with a specific eye on speech norms. I introduce Odessa – a DEcentralized Social Systems App at MIT.
- Dissertation title: Bridging communities: fostering pro-social networks with a focus on speech norms.
- Dissertation committee: Prof. Deb Roy, Prof. Jonathan Zittrain, Prof. Rosalind Picard.
- General Exam committee: Prof. Deb Roy, Prof. Veronica Barassi, Prof. Sasha Rush.
- Relevant classes: Machine Learning for Natural Language by Sasha Rush at Harvard University, Advances in Computer Vision by B. Freeman, A. Torralba, and P. Isola; Affective Computing by RW Picard.
- Awarded 2023 - [MIT Women of Excellence Award](#), MIT, USA.

Pontificia Universidad Católica de Chile (PUC) – Santiago, Chile

M.S. Machine Learning. GPA 6.7/7.0.

August 2016 – August 2017

B.S. Industrial Engineering Information Technology. GPA: 6.1/7.0.

March 2012 – August 2017

- Graduated with maximum distinction and ranked top 1% for both degrees.
- Awards and recognitions include:
 - Best Computer Science Thesis Award (2017).
 - Recipient of Academic Excellence Scholarship (2016).
 - Top National Score in Mathematics, Ministry of Education of Chile, University Admission (2012).
 - Recipient of Honor Scholarship (2012).
- Leader and co-founder, Plan Deportivo Ingeniería (2013 - 2017) to create, promote, and develop soccer, basketball, and volleyball teams in the school. Attracted around 500 students every semester (ongoing with new leaders).

RESEARCH AND PROFESSIONAL EXPERIENCE

Massachusetts Institute of Technology – Cambridge, MA, USA

Ph.D. Research Assistant, MIT Media Lab, Lab for Social Machines

September 2018 – present

- Developing and deploying community-aware machine learning models and systems.
- Led the process to architect, design, implement, and test [Odessa](#), a new decentralized social network app available for iOS and Android. Implemented state-of-the-art LLM-based algorithms to re-imagine new pro-social opportunities for healthier online spaces. This is being adopted across multiple institutions, such as Harvard Berkman Klein Center.
- Architected, designed, and deployed a randomized controlled experiment to help increase Twitter users' emotional awareness.
- Designed, developed, deployed, and data analyzed an in-the-wild study over four months with middle-schoolers attending a high-poverty school largely serving Black and Hispanic students to explore how a technology-mediated program might help bridge mentorship gaps. I was the leader in all the technology aspects.
- Developed and evaluated an approach to understand how different attention types relate to different types of questions in Visual Question Answering (VQA). [repo](#).

Google LLC – Remote, MA, USA

Research Intern, Google Research, Montreal Translate Team

June 2021 – August 2021

Manager: Research Scientist, George Foster, Ph.D.

- Worked at the Google Translate Research Group, which focuses on research and product.
- Rapidly learned Google’s development infrastructure to effectively drive the design, implementation, and experimentation that lead our team to draw conclusions and define research steps.
- Published and presented a [paper](#) at a top International Natural Language conference.

Falabella Retail S.A. – Santiago, Chile

Lead Data Scientist and Corporate Web Intelligence Engineer

September 2017 – June 2018

- Leader of the first data science team. Led our team to outperform the user experience and findability at Falabella.com from a machine learning perspective. Analyzed customer behavior from more than 100 GB of e-commerce navigation data; results fed into better organization and prioritization of product mix. The CEO officially acknowledged our team due to the produced increase in the conversion rate of the e-commerce site.
- Recruited and interviewed data scientists for the company. It included various profiles: data wranglers, engineers, scientists with mathematical and algorithmic mindsets, and data science team leaders.
- Worked with cross-functional teams, including business managers and software developers, to better understand user behavior and enhance user experience and the performance of the search engine.
- More details and a LinkedIn recommendation at: <https://belencarolina.com/industry>.

Machine Learning Intern

August 2016 – December 2016

- Full-stack development of the first in-house recommender system. As I was the first person to propose such a project, my manager and the CEO gave me full flexibility to explore and innovate within Falabella. These five months were instrumental for the company to grow in machine learning capacities.

Harvard University – Cambridge, MA, USA

Research Assistant

March 2016 – August 2017

- Developed my master’s thesis, as a paid RA at the Institute of Applied Computational Science (IACS), focused on machine learning, specifically on approximate variational inference for crowdsourcing labels.
- Member of the Harvard-Chile Data Science School, advised by Prof. Pavlos Protopapas, Scientific Program Director at IACS.
- Recipient of IACS financial aid to pursue my research (stipend and flight tickets).

Freelancer Consultant, Customer Intelligence

October 2016 – July 2017

- Developed optimized data preprocessing and models through tens of GB of data.
- Defined a set of quick-win promotions to find the relationship between brands and type of products.
- Classified clusters of customers according to their buying behavior to give personalized service.

TEACHING EXPERIENCE

Massachusetts Institute of Technology – Cambridge, MA, USA

Ph.D. Teaching Assistant at the Department of Electrical Engineering and Computer Science (EECS)

Fall 2019

- Served as grad TA for graduate (6.862) machine learning and undergraduate (6.036).

- Guiding grad students in their applied machine learning projects and developing material for new problem sets related to deep learning. Worked under professors L. Kaelbling, T. Broderick, D. Boning, and P. Jaillet, J. Andreas, among others.

Ph.D. Teaching Assistant at the MIT Media Lab

Fall 2020

- Serving as Chief TA for Understanding Public Thought, which is jointly offered at UW-Madison and MIT.
- Our goal was to design and advance strategies for measuring public opinion, paying particular attention to the need to understand perspectives as well as gauge preferences.
- Led and developed a series of hand-on workshops and lectures on natural language processing methods to analyze large datasets of online data to surface and understand public thought.

Instructor

- June 2022 and May 2023, Lead the Machine Learning and Data Science workshop the [MIT Micromasters](#) in Statistics and Data Science in partnership with MIT Global Startup Labs.
- 2022 Winter, MIT Teaching + Learning Lab. Instructor for sessions at the Teaching Assistant week.

Kaufman Teaching Certificate, [program](#).

Fall 2020

Teaching Development Fellow, [program](#).

2021 – 2023

Pontificia Universidad Católica de Chile – Santiago, Chile

Adjunct Professor of Data Mining

August 2017 – December 2017

- Recipient award for Exceptional Quality as Lecturer and Maximum Student Satisfaction (2018).
- Only professor teaching Data Mining, a core required class for in data science and machine learning for master and PhD students and for undergrad students majoring in data science.
- Designed curriculum, created lectures, assessments, and feedback process (doubled class attendance from previous years). Managing 4 teaching assistants.
- Youngest professor in the Computer Science Department of the School of Engineering.

Adjunct Professor of Data Mining and Management

March 2018 – July 2018

- Designed curriculum, lectures, assessments, and feedback process for a master's class of 50 students.
- Award: Obtained maximum student satisfaction in both scientific and practical knowledge acquired.
- Offered versions of this class a few times since then, including lectures to peers and managers at Falabella and as an invited lecturer to PUC classes at the CS department.

Adjunct Professor of Advanced Programming

March 2018 – July 2018

- Designed curriculum, lectures, assessments, and feedback process for an audience of 50 students.
- Core CS class for 2nd- and 6th-year (undergrad) university students.

Teaching Assistant at the Computer Science Department

2013 - 2017

- Served as Chief Teaching Assistant for 8 computer science and engineering classes (2013 – 2017), including Advanced Programming, Stochastic Models, Bayesian Inference, Data Structure and Algorithms, among others.
- Managed teams of 3-30 teaching assistants (academic and administrative management).
- Award: Recipient of the Exceptional Quality as Teaching Assistant Award (2015).

ACADEMIC PUBLICATIONS

- Overney, C., **Saldias, B.**, Dimitrakopoulou, D., & Roy, D. (March 2024) **SenseMate: An Accessible and Beginner-Friendly Human-AI Platform for Qualitative Data Analysis**. Proceedings of the 29th International Conference on Intelligent User Interfaces (IUI), Greenville, South Carolina. Contribution: Led the implementation, training setup, and evaluation framework for rationale extraction models. [paper](#).
- **Saldias, B.**, Foster, G., Freitag, M., & Tan, Q. (2022, May) Toward More Effective Human Evaluation for Machine Translation. *Proceedings of the 2022 ACL Workshop on Human Evaluation of NLP Systems (HumEval)*. ACL, Dublin, Ireland. [paper](#).
- **Saldias, B.**, & Roy, D. (2020, July) Exploring aspects of similarity between spoken personal narratives by disentangling them into narrative clause types. *Proceedings of the 2020 ACL Workshop on Narrative Understanding, Storylines, and Events (NUSE)*. ACL. [paper](#). [video](#). [repo](#).
- **Saldias, B.**, & Picard, R. (2019, September). Tweet Moodifier: Towards Giving Emotional Awareness to Twitter Users. *Proceedings of the 2019 International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE. Cambridge, UK. [paper](#). [media](#).
- **Saldias, B.**, Protopapas, P., & Pichara, K. (2019, May). A Full Probabilistic Model for Yes/No Type Crowdsourcing in Multi-Class Classification. *Proceedings of the 2019 SIAM International Conference on Data Mining (SDM)* (pp. 756-764). SIAM. Calgary, Canada. [paper](#).
 - Work was supported in part by the CSS survey, which is funded by the National Aeronautics and Space Administration under Grant No. NNG05GF22G issued through the Science Mission Directorate Near-Earth Objects Observations Program.
- Bhargava, R.*, Chung, A., Gaikwad, N., Hope, A., Jen, D., Rubinovitz, J., **Saldias, B.***, & Zuckerman, E. (2019, November). Gobo: A System for Exploring User Control of Invisible Algorithms in Social Media. *Submitted to the 2019 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)*. ACM. *All authors contributed equally to this research. [paper](#).
- **Saldías, Belén**. (2017, August) A full probabilistic model for yes/no type crowdsourcing in multi-class classification. *Santiago, Chile: Pontificia Universidad Católica de Chile*. (master's thesis). [paper](#). [repo](#). Advisors: Prof. Karim Pichara (PUC, Chile), and Prof. Pavlos Protopapas (Harvard, USA).

UNDER REVIEW

- **Saldias, B.***, Gillani, N.*, Makini, S.*, Hughes, M.*, & Roy, D.* (2021) Exploring a Human-Bot Partnership to Bridge Childhood Mentorship Gaps through Personal Narratives. [preprint](#). [web page](#). * Indicates contributions of equal importance.

OTHER ACADEMIC AND RELATED VENUES

- Posters:
 - **Saldias, B.**, (July 2024) Designing Child-Centered Content Exposure and Moderation. Presented as a poster at the 10th International Conference on Computational Social Science (IC2S2), Philadelphia, USA. [paper](#).
 - **Saldias, B.**, & Roy, D. (2020, July) Exploring aspects of similarity between spoken personal narratives by disentangling them into narrative clause types. WiML at ICML.
 - **Saldias, B.**, & Roy, D. (2019, December). Understanding story similarity for personal narratives. 14th Women in Machine Learning (WiML) Workshop co-located with NeurIPS. Vancouver, British Columbia, Canada. Awarded a travel grant, to attend the workshop, by WiML.

- Book contributions:
 - Pichara, K., & Pieringer, C., (2017). *Advanced Computer Programming in Python*. I served as the main contributor of this published book as I led developing and testing the code exercises and examples in the book. Freely available at <https://advancedpythonprogramming.github.io> (also [in Amazon](#)).
- Leader at workshop sessions:
 - **Saldias, B.**, & Ali, S. (2020, July). Abstract: Towards Children-Aware Machine Learning with a Focus on NLP Challenges and Applications. 1st WiML Un-Workshop co-located with ICML. [repo](#).
- Noteworthy acknowledgments:
 - Hughes, M., & Roy, D. (CHI 2021) *Keeper: A Synchronous Online Conversation Environment Informed by In-Person Facilitation Practices*. [paper](#). [thesis](#). Contribution: I architected and developed the complete backend, server setup, and deployment processes to make Keeper functional and allow for human evaluation (as also acknowledged in the MIT M.S. Thesis linked above.).
 - Jiang, H., et al. (COLING 2022) *CommunityLM: Probing Partisan Worldviews from Language Models*. [paper](#). Contribution: Discussion in ideation, framing, and feedback stages.
- Related grants:
 - 2022 - Travel Grants, MIT Media Lab (\$2K).
 - 2021 - 2023 MIT Teaching Development Fellowship (\$4K).
 - 2019 - Travel Grants, MIT Media Lab (\$6K) and WiML (\$1K).

LEADERSHIP AND SERVICE

- 2023, General Chair and Keynote, MIT-Chile Research Workshops in Human-Centered Artificial Intelligence and Data Visualization.
 - Chaired and organizer first workshop of its kind in Chile by leading a team of 3 program chairs. Hosted two locations Santiago and Concepción.
 - Individually raised 30K from prestigious sponsors, including MIT, Pontificia Universidad Católica de Chile, and multiple nation-wide research institutions.
 - Invited and hosted 20 speakers (flying from multiple top universities, including MIT, Harvard, and Notre Dame) and 8 lecturers with topics ranging from bias in machine learning, to brain-computer interaction, and activities from drawing and storytelling with data to a live EEG demo. [See more](#).
 - Hosted more than 100 students, out of which 44% self-identify as women or non-binary, and 25% are not from Santiago (traveling internationally and from northern and southern regions of Chile—for whom we provided travel funding).
 - Grants received by: MIT Global Experiences, MIT Media Lab, MIT Center for Constructive Communication, MIT Teaching Lab, Harvard Institute for Applied Computational Science, Chile's Millennium Institute on Intelligent Healthcare Engineering, Chile's Millennium Institute Foundational Research on Data, Chile's National Center for Artificial Intelligence Research, Chile's Millennium Nucleus on Young Exoplanets and their Moons, Amazon Web Services, University of Concepción, Pontificia Universidad Católica de Chile.
- 2020 – 2021, Executive Council Co-Chair, MIT Sidney Pacific Graduate Community. MIT Sidney Pacific homes nearly 650 graduate students from many corners of the world.

- Service to the academic community includes:
 - Program Committee at [FAcct2023](#), 2023.
 - Program Committee at Data-driven Humanitarian Mapping, [KDD](#), 2022.
 - Program Committee and Volunteer at [MIT Path to Professorship](#), 2022.
 - Board of Directors at Women in Machine Learning (WiML), 2022 – 2023.
 - Reviewer for CHI 2021 Papers.
 - Logistics and Funding Chair for the Workshop for Women in Machine Learning at NeurIPS 2020.
 - Super-volunteer for the 1st WiML Un-Workshop at ICML, 2020.
 - Reviewer for the ICML Deploy and Monitor ML workshop, 2020.
 - Reviewer for the 14th NeurIPS WiML workshop, 2019.

- Invited lectures and presentations include:
 - TEDx Speaker, invited by [TEDx Bentley U](#). April 2023.
 - Guest Lecture, Korea Advanced Institute of Science and Technology, hosted by Prof. Hyeon Yeo. May 2024.
 - Guest Lecture, MIT, hosted by Prof. Roz Picard. Title: LLMs, responsible sentiment analysis, and content moderation in social media. October 2023.
 - Guest Lecture at [Ars Electronica](#), invited by Johannes Kepler University Linz, Austria, Sept 2023.
 - Guest Lecture, MIT MISTI and MIT GSL. Title: Intro to Machine Learning. May 2023 and June 2022.
 - Guest Lecture, MIT, hosted by Prof. Roz Picard. Title: Responsible sentiment analysis, affect manipulation and social media. October 2022.
 - Guest Lecture, Research in Machine Learning at PUC CS/DCC Chile, PhD Seminar, October 2020.
 - Guest Lecture, PUC Chile. Title: Probabilistic Models for Clustering, April 2022.
 - Research Presentation, MIT, hosted by Prof. Joseph Paradiso. October 2023.
 - Research Presentation, MIT, hosted by Generative AI for Constructive Communication, MIT. March 2023.
 - Research Presentation, MIT, hosted by Prof. Roz Picard, MIT, Spring 2019, 2021.
 - Research Presentation, PUC Chile, hosted by Prof. Denis Parra at the Social Computing & Visualization Group. October 2019 and 2020.
 - Keynote Talk and Lecture, MIT – Chile Research Workshop, January 2023.
 - Keynote speaker, hosted by Agrosuper, Chile. Title: Machine Learning in Industry. August 2022.
 - Keynote speaker, Data Science in Industry at U. of Concepción, Chile, August 2020.
 - Keynote speaker, Innovation and Entrepreneurship, U. La Sabana, Colombia, October 2020.

- Discussion panels:
 - Developing a Research Thesis (2021, October), Introduction to Media Arts and Sciences at MIT.
 - Mitigating Bias in Analytics (2020, August). Hosted by Women in Analytics and Tableau [video](#).
 - Women in Data Science (2018, March). Santiago, Chile. Spoke to how to push data science and machine learning in the Chilean industry. [website](#).

- Research mentorship for multiple MIT master students and undergrad researchers, included as a crucial mentor and collaborator in their thesis.

PRINTED MEDIA

- 2024, Berlin, Germany. [Featured](#) in exhibition on art, AI, & work at Museum for Communication Berlin.
- 2023, Chile. “Chileans in Artificial Intelligence.” Diario Financiero, DF MAS. Featured in the most important financial and economic portal in Chile. [online](#). [pdf](#).

- 2020, Chile. “Six women leave footprints in Boston, one of the most important tech hubs in the world.” El Mercurio, Chile, 2020. Featured in El Mercurio's cover and in the Innovation section's cover. El Mercurio is considered Chile's newspaper of record, and it is considered the oldest daily in the Spanish language currently in circulation. [pdf](#).
- 2019, Spain. “This is how emotions spread on Twitter.” El País, España. Article based on Saldias, B., & Picard, R. (2019, September). El País is a Spanish-language daily newspaper in Spain. According to the Office of Justification of Dissemination, it is the second most circulated daily newspaper in Spain as of December 2017. [link](#).
- 2019, Spain. “Research in Spanish at the MIT Media Lab.” La Razón, España. Article highlighting spanish-speaking researchers at the MIT Media Lab. La Razón is a daily newspaper based in Madrid, Spain. It has the sixth-highest circulation among general-interest Spanish dailies. [link](#).
- 2023, Chile. “Chileans at MIT.” Featured in El Mercurio for my work in artificial intelligence. [online](#). [pdf](#).

ONLINE MEDIA

- 2021, Cambridge, MA, USA. “Reinventing the graduate community during the pandemic at MIT.” Featured at MIT News for my contribution to the community. [online](#).
- 2023, USA. TEDx talk featured on Journal of the Data Visualization Society (DVS). [online](#).

OTHER

- Languages: Spanish (native speaker), English (full professional working proficiency).