

Max Bartolo

Education

- since Jan 2019 **PhD in Computer Science and Artificial Intelligence**
University College London, UK.
I am a final-year PhD student at the UCL Natural Language Processing (NLP) group under the supervision of Pontus Stenetorp and Sebastian Riedel. My current research revolves around understanding and improving the capabilities and robustness of contemporary NLP models.
- Sep 2016 – Sep 2017 **MSc Business Analytics (Computer Science specialisation)**
University College London, UK, Distinction, Dean's List.
Courses followed included Statistical Natural Language Processing, Introduction to Supervised Learning, Business Strategy & Analytics, Graphical Models, Programming, Research Software Engineering with Python, Applied Machine Learning, Web Economics, Data Science for Spatial Systems, Data Analytics and Advanced Topics in Machine Learning (delivered by Google Deepmind).
- Sep 2008 – Jul 2012 **BEng Mechanical Engineering (Structural Mechanics and Thermofluids)**
University of Malta, Malta, 2:1 Honours.

Experience

- May 2022 – Sep 2022 **Research Scientist Intern, DeepMind, London.**
Research on large language models under the guidance of Po-Sen Huang and Johannes Welbl with the Robust and Verified AI (RVAI) team.
- Apr 2021 – Jul 2021 **External Research Collaborator, Facebook AI Research (FAIR), London.**
Research on dynamic adversarial benchmarking and evaluation, visual grounding, and generative assistive models for more effective adversarial annotation. Worked under the guidance of Douwe Kiela and Robin Jia with publications at NAACL '22, ACL '22 and CVPR '22.
- Sep 2020 – Dec 2020 **Research Scientist Intern, Facebook AI Research (FAIR), London.**
Research revolving around understanding how NLP models behave in the face of adversarial inputs, improving model robustness, and challenging models beyond their current capabilities through human-and-machine generative processes. Worked under the guidance of Douwe Kiela and Robin Jia with publications at NAACL '21 and EMNLP '21.
- Feb 2020 – Sep 2020 **ConceptionX, Cohort III, London.**
Formed part of Cohort III of the ConceptionX start-up accelerator programme.
- since Mar 2020 **Adjunct Teaching Fellow, University College London, UK.**
Developed and deliver the MSIN0221 Natural Language Processing (NLP) module at the UCL SoM with excellent student feedback. I also supervise and support student dissertations in close collaboration with industry.
- since Feb 2019 **Bartolo AI, Founder & CEO, London.**
Working with companies ranging from startups to multinationals on their AI implementation strategies and real-world application of cutting-edge Natural Language Processing (NLP) research.

- May 2017 – **Machine Learning Engineer**, *Bloomsbury AI*, London.
- Sep 2018 Played a significant role in developing a large-scale open-domain question answering system, managing client relationships, building a highly-skilled technical team and guiding company direction towards a successful exit. Co-authored a research paper proposing a novel task definition, dataset, and baseline models on *Conversational Machine Reading* (EMNLP '18).
- Feb 2017 – **Data Scientist**, *Satalia*, London.
- May 2017 Prototyped a pay structure recommender pipeline based on the analysis of key network metrics between people and skills within the company, primarily using NLP techniques such as sentiment analysis and topic modelling from intra-organisation communication channels.
- May 2014 – **CyberMax Creations**, *Founder & CEO*, Malta.
- Sep 2018 Web development, design and digital marketing as well as sub-contracted work with local agencies.
- May 2014 – **Coach & Co-founder**, *Swieqi United FC*, Malta.
- Aug 2017 Co-founder, coach and development manager of the Swieqi United Women's Football Team.
- Aug 2013 – **Operations Manager**, *Step Enterprises Ltd*, Malta.
- May 2014 Implemented a new sales structure and broadened service offering. Actively worked with business leaders to improve operational efficiency and implement ISO 9001 Quality Management Systems.
- Jul 2011 – **Quality Engineer**, *Methode Electronics*, Malta.
- Aug 2013 Responsible for problem analysis, root cause identification and resolution in liason with OEM customers including Audi, BMW, Ford, Jaguar Land Rover, McLaren, Mitsubishi and others. Appointed Quality Leader on two project launches, supporting a team of quality, design and manufacturing engineers from project conception.

Master Thesis

- Title *Transfer Learning for Open-Domain Question Answering*
- Supervisors Sebastian Riedel, Tim Rocktäschel, Guillaume Bouchard
- Description Demonstrated that learning acquired from Wikipedia articles could be effectively transferred to the legal, travel and open-domain settings with relatively little additional data. Developed a question answering pipeline involving context extraction, data annotation and validation, context retrieval and question answering focused on machine comprehension, and integrated this into the *Bloomsbury AI* product.

Invited Talks

- 2022 **Dynamic Adversarial Data Collection**, *London*, UK.
Invited talk at the *Oracle Labs ML Seminar Series*.
- Dynamic Adversarial Data Collection for Question Answering**, *London*, UK.
Invited talk at the *University of Oxford*.
- 2021 **Dynamic Benchmarking**, *London*, UK.
Invited talk at the *UCL AI Centre 2nd Anniversary Showcase*.
- Adversarial Human Annotation and Dynamic Benchmarking**, *London*, UK.
Guest Lecture at UCL for the COMP0090 Introduction to Deep Learning module.
- 2020 **Generative Data Augmentation for Improved QA Model Robustness**, *London*, UK.
Presented at the *Facebook AI Research NLP & Conversational AI Meeting*.
- Dynamic Benchmarking and Evaluation**, *London*, UK.
Invited talk at the *UCL AI Centre* session on AI in science, industry and society at *TheAlgo2020*.
- Humans-and-Machines in the Loop for Dynamic Benchmarking**, *London*, UK.
Invited talk at the Annual *MURI* Review Meeting.

Adversarial Human Annotation for Reading Comprehension, *London, UK.*

Invited talk at the *University of Cambridge NLIP Seminar Series*.

2019 **Asking Harder Questions**, *London, UK.*

Invited talk at the *UCL NLP Inaugural Event*.

Question Answering using Rules and Conversation, *London, UK.*

Invited talk at the *South England Natural Language Processing meetup*.

2017 **Question Answering at Scale**, *London, UK.*

Invited talk at the "A Common Language for Intelligence" meetup.

Teaching

since 2020 **Module Lead**, *University College London, UK.*

Developed and deliver the MSIN0221 NLP module at the UCL SoM.

2019 **Guest Lecturer**, *University College London, UK.*

Overview of NLP (two lectures)

Guest Lecturer, *Peking University HSBC Business School (PHBS), Oxford, UK.*

Introduction to Python and Machine Learning Workshop

Teaching Assistant, *University College London, UK.*

COMP0090: Introduction to Deep Learning

Supervision

2021 **Supervisor**, *University College London, UK.*

Steven George, MSc Machine Learning

Supervisor, *University College London, UK.*

Oliver Lockett, MSc Business Analytics

Supervisor, *University College London, UK.*

Frederick Reid, MSc Business Analytics

Supervisor, *University College London, UK.*

Jake Rutherford, MSc Business Analytics

Co-supervisor, *University College London, UK.*

Benjamin Sidi, MSc Machine Learning

2020 **Supervisor**, *University College London, UK.*

Sam Hosegood, MSc Computational Statistics and Machine Learning

Investigating Manually Generated Adversarial Attacks on Transformer Models for Question Answering

2019 **Co-supervisor**, *University College London, UK.*

Louise Davies, MSc Computer Science

Towards Explainability by Clustering Hidden Representations

Community

Organisation

First Workshop on Dynamic Adversarial Data Collection (DADC) at NAACL'22 (<https://dadcworkshop.github.io/>).

HAMLETS (Human And Machine in-the-Loop Evaluation and Learning Strategies) Workshop at NeurIPS 2020 (<https://hamlets-workshop.github.io/>).

Reviewing

ACL Rolling Review (ARR), North American Chapter of the Association for Computational Linguistics (NAACL), Annual Meeting of the Association for Computational Linguistics (ACL), Conference on Neural Information Processing Systems (NeurIPS), Conference on Empirical Methods in NLP (EMNLP), Workshop on Representation Learning for NLP (Repl4NLP), Workshop on Machine Reading for Question Answering (MRQA), Conference on Automated Knowledge Base Construction (AKBC), the International Conference on Computational Linguistics (COLING).

Open Source Software

- Dynabench A research platform for dynamic data collection and benchmarking designed to address well-known issues with static benchmarks such as saturation, are susceptibility to overfitting, annotation artefacts and unclear or imperfect evaluation metrics.
- Cape A large-scale open-domain question answering system backed by an easy-to-use API giving users answers to questions about the contents of their documents as well as a means to correct, calibrate and re-train the models.

Certifications & Achievements

- 2022 Outstanding Paper Award for “Fantastically Ordered Prompts and Where to Find Them: Overcoming Few-Shot Prompt Order Sensitivity”, *ACL*.
- 2017 Dean’s List Award for Outstanding Academic Achievement, *UCL*.
- 2016 UEFA ‘C’ Coaching License, *Malta Football Association*.
- 2016 Google Analytics IQ, *Google*.
- 2015 Advanced Open Water Scuba Diver license, *PADI*.
- 2015 First Aid, *St John’s Ambulance*.
- 2015 The Data Scientist’s Toolbox, *Coursera C49ZCNDX5Y*.
- 2014 Nautical License, *Transport Malta*.
- 2006 1st Dan Shotokan Karate Black Belt, *Malta Karate Federation*.

Publications

- [1] **Max Bartolo**, Tristan Thrush, Sebastian Riedel, Pontus Stenetorp, Robin Jia, and Douwe Kiela. Models in the loop: Aiding crowdworkers with generative annotation assistants, 2021.
- [2] Tristan Thrush, Ryan Jiang, **Max Bartolo**, Amanpreet Singh, Adina Williams, Douwe Kiela, and Candace Ross. Winoground: Probing vision and language models for visio-linguistic compositionality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5238–5248, June 2022.
- [3] Tristan Thrush, Kushal Tirumala, Anmol Gupta, **Max Bartolo**, Pedro Rodriguez, Tariq Kane, William Gaviria Rojas, Peter Mattson, Adina Williams, and Douwe Kiela. Dynatask: A framework for creating dynamic AI benchmark tasks. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 174–181, Dublin, Ireland, May 2022. Association for Computational Linguistics.
- [4] **Max Bartolo**, Tristan Thrush, Robin Jia, Sebastian Riedel, Pontus Stenetorp, and Douwe Kiela. Improving Question Answering Model Robustness with Synthetic Adversarial Data Generation. In *EMNLP (Main Conference)*, 2021.

- [5] Maximilian Mozes, **Max Bartolo**, Pontus Stenetorp, Bennett Kleinberg, and Lewis D. Griffin. Contrasting Human- and Machine-Generated Word-Level Adversarial Examples for Text Classification. *In EMNLP (Main Conference)*, 2021.
- [6] Yao Lu, **Max Bartolo**, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity (**outstanding paper award**). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8086–8098, Dublin, Ireland, May 2022. Association for Computational Linguistics.
- [7] Douwe Kiela, **Max Bartolo**, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts, and Adina Williams. Dynabench: Rethinking benchmarking in NLP. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4110–4124, Online, June 2021. Association for Computational Linguistics.
- [8] Daniel Hoadley, **Max Bartolo**, Robin Chesterman, A Faus, W Hernandez, B Kultys, AP Moore, E Nemsic, N Roche, Jiangbo Shangguan, et al. A Global Community of Courts? Modelling the use of Persuasive Authority as a Complex Network. *Frontiers in Physics*, 9:331, 2021.
- [9] **Max Bartolo**, Alastair Roberts, Johannes Welbl, Sebastian Riedel, and Pontus Stenetorp. Beat the AI: Investigating Adversarial Human Annotation for Reading Comprehension. In *Transactions of the Association for Computational Linguistics (TACL)*. Presented at EMNLP, 2020.
- [10] Johannes Welbl, Pasquale Minervini, **Max Bartolo**, Pontus Stenetorp, and Sebastian Riedel. Undersensitivity in Neural Reading Comprehension. *Findings of EMNLP*, 2020.
- [11] Marzieh Saeidi*, **Max Bartolo***, Patrick Lewis*, Sameer Singh, Tim Rocktäschel, Mike Sheldon, Guillaume Bouchard, and Sebastian Riedel. Interpretation of Natural Language Rules in Conversational Machine Reading. *Proceedings of EMNLP*, 2018.