

Bloomberg AI

Bloomberg is building the world's most trusted information network for financial professionals. Our core product, the Bloomberg Terminal, is an independent and unbiased source of information for our clients – everyone from C-Suite executives, traders, analysts, government officials, business heavyweights, and news professionals around the globe.

6,000+ technologists – and counting – are dedicated to advancing and building new solutions for the Bloomberg Terminal in order to solve complex, real-world problems. No other company processes the breadth and depth of financial data into meaningful and actionable information as well, or as quickly, as we do. Our engineers and researchers work together to roll out new software daily that integrates seamlessly and immediately into our clients' workflows.

Our technology expertise is changing the world and makes a positive impact in our local communities. Bloomberg engineers are active members of the open source community, and contribute to and leverage open source software in our products. And, nearly all of our company profits go directly to Bloomberg Philanthropies to support lasting change in the environment, government innovation, education, arts, and public health.

Artificial Intelligence @ Bloomberg

We are a close-knit team of 150+ researchers and engineers, focusing on projects related to AI, ML, NLP, NLU, IR, and QA.

We collaborate with researchers from academia by publishing papers, providing funding for research and fellowships, participating in conferences and program committees, hosting interns, and peer-reviewing scientific articles.

We develop solutions for semantic and syntactic parsing problems, NER & NEL systems, sentiment analysis, market impact indicators, knowledge base construction, social media analysis, question answering systems, and high performance IE. Our customers rely on these systems to make swift financial decisions; therefore, we need to guarantee precision, accuracy and latency beyond what academia and many other industries would demand.

We are looking for research scientists and engineers who are eager to apply Machine Learning and NLP to solve challenging open-ended problems. If you want to be part of a team making an impact on the global financial industry and are not afraid to get your hands dirty in data, check out the following opportunities within our Artificial Intelligence group:

New York: https://bit.ly/ACL20_NYC

London: https://bit.ly/ACL20_LON (full-time)

https://bit.ly/ACL20_LON_DSI (internship)

Engineering

Bloomberg

Bloomberg @ ACL

Papers & System Demo

Analyzing Political Parody in Social Media

Antonis Maronikolakis, Danae Sánchez-Villegas, **Daniel Preoțiuc-Pietro** and Nikolaos Aletras
Session 8A Computational Social Science & Social Media-6 | Tuesday, July 7, 2020 | 8-9 AM EDT
Session 9B Computational Social Science & Social Media-7 | Tuesday, July 7, 2020 | 2-3 PM EDT

The Summary Loop: Learning to Write Abstractive Summaries Without Examples

Philippe Laban, **Andrew Hsi**, John Canny, and Marti Hearst
Session 9A Summarization-4 | Tuesday, July 7, 2020 | 1:00-2:00 PM EDT
Session 10B Summarization-5 | Tuesday, July 7, 2020 | 5:00-6:00 PM EDT

NSTM: Real-Time Query-Driven News Overview Composition at Bloomberg (System Demo)

Joshua Bambrick, Minjie Xu, Guim Perarnau, Igor Malioutov, Andy Almonte, Vittorio Selo and Iat Chong Chan

Demo Session 3B-3 | Wednesday, July 8, 2020 | 8:45-9:45 AM EDT
Demo Session 5B-3 | Wednesday, July 8, 2020 | 4:45-5:45 PM EDT

Temporally-informed Analysis of Named Entity Recognition

Shruti Rijhwani and **Daniel Preoțiuc-Pietro**
Session 13B Information Extraction-8 | Wednesday, July 8, 2020 | 9:00-10:00 AM EDT
Session 15A Information Extraction-11 | Wednesday, July 8, 2020 | 4:00-5:00 PM EDT

Multi-Domain Named Entity Recognition with Genre-Aware and Agnostic Inference

Jing Wang, Mayank Kulkarni and **Daniel Preoțiuc-Pietro**
Session 14B Information Extraction-10 | Wednesday, July 8, 2020 | 2:00-3:00 PM EDT
Session 15B Information Extraction-12 | Wednesday, July 8, 2020 | 5:00-6:00 PM EDT

ECNLP 3 Workshop | Friday, July 10, 2020

<https://sites.google.com/view/ecnlp/acl-2020>

Semi-Supervised Iterative Approach for Domain-Specific Complaint Detection in Social Media

Akash Gautam, **Rakesh Gosangi**, Debanjan Mahata, Rajiv Ratn Shah

Read more about the work of Bloomberg's engineers and data scientists and view a selection of their published research on AI, Machine Learning and NLP: [TechAtBloomberg.com/AI](https://techatbloomberg.com/AI).

Learn more

// Email our recruiters to learn more about full-time, co-op, post-doc, and internship roles:

New York: Irinka - itoidze@bloomberg.net

London: Fran - fcucinotta@bloomberg.net

// Learn about our data science research grants & fellowships:

[TechAtBloomberg.com/data-science-research-grant-program/](https://techatbloomberg.com/data-science-research-grant-program/)

// Explore our Foundations of Machine Learning video course:

[TechAtBloomberg.com/FOML](https://techatbloomberg.com/FOML)

// Read about the unique problems we're solving:

[TechAtBloomberg.com/post-topic/data-science](https://techatbloomberg.com/post-topic/data-science)

// Explore our open source projects: www.github.com/Bloomberg

// Follow us on Twitter: [@TechAtBloomberg](https://twitter.com/TechAtBloomberg)

TechAtBloomberg.com