Language at DeepMind

Language is at the core of DeepMind's research. The Language Team studies all aspects of language processing both in artificial agents and in humans. We take a multi-disciplinary approach to studying language combining insights from machine learning, linguistics, cognitive science, and computer vision. Our group works on a broad range of topics such as:

- Models of language acquisition and processing.
- Language tasks and applications such as translation, question answering, and dialogue.
- Emergence of communication protocols in multi-agent frameworks.
- Multimodal learning and its applications.
- NLP for low-resource languages.
- Analysing and explaining neural models.
- Learning useful general-purpose representations.
- Leveraging linguistic structure in neural models.

ACL 2020 Papers

- **Multi-agent Communication meets Natural Language: Synergies between Functional and Structural Language Learning**
  Angeliki Lazaridou, Anna Potapenko, and Olivier Tieleman access here
- **Learning to Segment Actions from Observation and Narration***
  Daniel Fried*, Jean-Baptiste Alayrac, Phil Blunsom, Chris Dyer, Stephen Clark, and Aida Nematzadeh access here
- **Do Transformers Need Deep Long-Range Memory?**
  Jack Rae access here
- **On the Cross-lingual Transferability of Monolingual Representations***
  Mikel Artetxe*, Sebastian Ruder, and Dani Yogatama access here
- **A Call for More Rigor in Unsupervised Cross-lingual Learning***
  Mikel Artetxe*, Sebastian Ruder, Dani Yogatama, Gorka Labaka, and Eneko Agirre access here
- **A Probabilistic Generative Model for Typographical Analysis of Early Modern Printing**
  Kartik Goyal, Chris Dyer, Christopher Warren, Maxwell G'Sell, and Taylor Berg-Kirkpatrick access here
- **Make Up Your Mind! Adversarial Generation of Inconsistent Natural Language Explanations**
  Oana-Maria Camburu, Brendan Shillingford, Pasquale Minervini, Thomas Lukasiewicz, and Phil Blunsom access here

*Papers co-authored by DeepMind interns

ACL 2020 Schedule

Main Conference

Monday July 6th

Slot 1 - 18:00| Slot 2 - 21:00 UTC+0

Learning to Segment Action from Observation and Narration

Daniel Fried (UC Berkeley), Jean-Baptiste Alayrac, Phil Blunsom, Chris Dyer, Stephen Clark, and Aida Nematzadeh

Tuesday July 7th

Slot 1 - 08:00 | Slot 2 - 13:00 UTC+0

Better Document-level Machine Translation with Bayes; Rule

Lei Yu, Laurent Sartran, Wojciech Stokowiec, Wang Ling, Lingpeng Kong, Chris Dyer and Phil Blunsom

Slot 1 - 13:00 | Slot 2 - 18:00 UTC + 0

On the Cross-lingual Transferability of Monolingual Representation

Mikel Artetxe (University of Basque Country), Sebastian Ruder and Dani Yogatama

Wednesday July 8th

Slot 1 - 9:00 | Slot 2 - 13:00 UTC + 0

A Call for More Rigor in Unsupervised Cross-lingual Learning

Mikel Artetxe (University of Basque Country)*, Sebastian Ruder*, Dani Yogatama, Gorka Labaka (University of Basque Country), and Eneko Agirre (University of Basque Country)

Slot 1 - 13:00 | Slot 2 - 17:00 UTC + 0

Multi-agent Communication Meets Natural Language: Synergies between Functional and Structural Language Learning

Angeliki Lazaridou, Anna Potapenko and Olivier Tieleman

Slot 1 - 12:00 | Slot 2 - 18:00 UTC + 0

Do Transformers Need Deep Long-Range Memory?

Jack Rae and Ali Razavi