

NAACL HLT 2019

**NeuralGen Workshop: Methods for Optimizing and  
Evaluating Neural Language Generation**

**Proceedings of the First Workshop**

June 6, 2019  
Minneapolis, Minnesota

©2019 The Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)  
209 N. Eighth Street  
Stroudsburg, PA 18360  
USA  
Tel: +1-570-476-8006  
Fax: +1-570-476-0860  
[acl@aclweb.org](mailto:acl@aclweb.org)

ISBN 978-1-950737-02-4

## Introduction

NeuralGen is the first workshop on Methods for Optimizing and Evaluating Neural Language Generation, being held at NAACL 2019 in Minneapolis, Minnesota. The goal of this workshop is to discuss new frontiers for language generation that address some of the recurring problems in existing techniques (eg. bland, repetitive language). More specifically, this workshop is aimed at sharing novel modeling techniques that go beyond maximum likelihood training, new techniques for robust evaluation and interpretation of model output, and strategies for generalization of generation systems.

We are pleased to have received 42 submissions, covering a wide range of topics related to modeling, evaluation and analysis of novel generation systems. 17 of the submissions have been accepted into the final program (approximately 40% acceptance rate). The workshop schedule includes 11 archival papers and 17 poster presentations. We are also thankful to have seven invited speakers: Kyunghyun Cho, He He, Graham Neubig, Yejin Choi, Alexander Rush, Tatsunori Hashimoto, and Hal Daumé III. The workshop also includes a panel discussion from the speakers and spotlight talks for a selection of accepted papers.

We would like to thank our invited speakers, authors, and reviewers for contributing to our program. Additionally, we would like to express gratitude to our sponsors, who have been generous in supporting the workshop.

*Antoine Bosselut, Asli Celikyilmaz, Srinivasan Iyer, Marjan Ghazvininejad, Urvashi Khandelwal, Hannah Rashkin, Thomas Wolf*



**Organizers:**

Antoine Bosselut, University of Washington  
Asli Celikyilmaz, Microsoft Research  
Srinivasan Iyer, University of Washington  
Marjan Ghazvininejad, Facebook AI Research  
Urvashi Khandelwal, Stanford University  
Hannah Rashkin, University of Washington  
Thomas Wolf, HuggingFace

**Steering Committee:**

Yejin Choi, University of Washington and Allen Institute for AI  
Dilek Hakkani-Tür, Amazon Research  
Dan Jurafsky, Stanford University  
Alexander Rush, Harvard University

**Program Committee:**

Adji Dieng, Columbia University  
Matthew Peters, Allen Institute for AI  
Alexander Rush, Harvard University  
Maxwell Forbes, University of Washington  
Andrew Hoang, University of Washington  
Mohit Iyyer, University of Massachusetts, Amherst  
Angela Fan, Facebook AI Research  
Nelson Liu, University of Washington  
Ari Holtzman, University of Washington  
Ofir Press, University of Washington  
Arun Chaganty, Eloquent Labs  
Ondrej Dusek, Heriot-Watt University  
Caiming Xiong, Salesforce Research  
Po-Sen Huang, Deepmind  
Chandra Bhagavatula, Allen Institute for AI  
Ramakanth Pasunuru, University of North Carolina  
Dallas Card, Carnegie Mellon University  
Rik Koncel-Kedziorski, University of Washington  
Daniel Fried, University of California, Berkeley  
Roei Aharoni, Bar-Ilan University  
Dinghan Shen, Duke University  
Ronan Le Bras, Allen Institute for AI  
Elizabeth Clark, University of Washington  
Saadia Gabriel, University of Washington  
Gabi Stanovsky, Allen Institute for AI  
Sachin Mehta, University of Washington  
He He, New York University  
Sam Wiseman, Toyota Technological Institute at Chicago

Ioannis Konstas, Heriot-Watt University  
Sebastian Gerhmann, Harvard University  
Jan Buys, University of Washington  
Siva Reddy, Stanford University  
Jesse Dodge, Carnegie Mellon University  
Spandana Gella, Amazon AI  
John Wieting, Carnegie Mellon University  
Tatsunori Hashimoto, Stanford University  
Keisuke Sakaguchi, Allen Institute for AI  
Terra Blevins, University of Washington  
Lianhui Qin, University of Washington  
Vered Shwartz, Bar-Ilan University  
Lifu Huang, Rensselaer Polytechnic Institute  
Verena Rieser, Heriot-Watt University  
Maarten Sap, University of Washington  
Yangfeng Ji, University of Virginia

**Invited Speakers:**

Kyunghyun Cho, New York University and Facebook AI Research  
Yejin Choi, University of Washington and Allen Institute for AI  
Hal Daumé III, University of Maryland and Microsoft Research  
Tatsunori Hashimoto, Stanford University  
He He, New York University and Amazon Web Services  
Graham Neubig, Carnegie Mellon University  
Alexander Rush, Harvard University

## Table of Contents

|  |    |
|--|----|
| <i>An Adversarial Learning Framework For A Persona-Based Multi-Turn Dialogue Model</i><br>Oluwatobi Olabiyi, Anish Khazane, Alan Salimov and Erik Mueller .....                                    | 1  |
| <i>DAL: Dual Adversarial Learning for Dialogue Generation</i><br>Shaobo Cui, Rongzhong Lian, Di Jiang, Yuanfeng Song, Siqi Bao and Yong Jiang.....   | 11 |
| <i>How to Compare Summarizers without Target Length? Pitfalls, Solutions and Re-Examination of the Neural Summarization Literature</i><br>Simeng Sun, Ori Shapira, Ido Dagan and Ani Nenkova ..... | 21 |
| <i>BERT has a Mouth, and It Must Speak: BERT as a Markov Random Field Language Model</i><br>Alex Wang and Kyunghyun Cho .....  | 30 |
| <i>Neural Text Simplification in Low-Resource Conditions Using Weak Supervision</i><br>Alessio Palmero Aprosio, Sara Tonelli, Marco Turchi, Matteo Negri and Mattia A. Di Gangi ...                | 37 |
| <i>Paraphrase Generation for Semi-Supervised Learning in NLU</i><br>Eunah Cho, He Xie and William M. Campbell .....  | 45 |
| <i>Bilingual-GAN: A Step Towards Parallel Text Generation</i><br>Ahmad Rashid, Alan Do Omri, Md Akmal Haidar, Qun Liu and Mehdi Rezagholizadeh .....   | 55 |
| <i>Designing a Symbolic Intermediate Representation for Neural Surface Realization</i><br>Henry Elder, Jennifer Foster, James Barry and Alexander O'Connor .....                                   | 65 |
| <i>Neural Text Style Transfer via Denoising and Reranking</i><br>Joseph Lee, Ziang Xie, Cindy Wang, Max Drach, Dan Jurafsky and Andrew Ng .....  | 74 |
| <i>Better Automatic Evaluation of Open-Domain Dialogue Systems with Contextualized Embeddings</i><br>Sarik Ghazarian, Johnny Wei, Aram Galstyan and Nanyun Peng .....                              | 82 |
| <i>Jointly Measuring Diversity and Quality in Text Generation Models</i><br>Danial Alihosseini, Ehsan Montahaei and Mahdieh Soleymani Baghshah .....   | 90 |





# Conference Program

Thursday, June 6, 2019

9:00–9:10 *Opening Remarks*

9:10–9:45 **Invited Talk: Kyunghyun Cho**

9:45–10:20 **Invited Talk: He He**

10:20–10:35 *Coffee Break*

10:35–11:10 **Invited Talk: Graham Neubig**

11:10–11:45 **Invited Talk: Yejin Choi**

11:45–13:15 *Lunch*

13:15–13:50 **Invited Talk: Alexander Rush**

13:50–14:15 **Spotlight Talks**

14:15–15:45 **Poster Session**

*An Adversarial Learning Framework For A Persona-Based Multi-Turn Dialogue Model*

Oluwatobi Olabiyi, Anish Khazane, Alan Salimov and Erik Mueller

*DAL: Dual Adversarial Learning for Dialogue Generation*

Shaobo Cui, Rongzhong Lian, Di Jiang, Yuanfeng Song, Siqi Bao and Yong Jiang

*Towards Coherent and Engaging Spoken Dialog Response Generation Using Automatic Conversation Evaluators*

Sanghyun Yi, Rahul Goel, Chandra Khatri, Tagyoung Chung, Behnam Hedayatnia, Anu Venkatesh, Raefer Gabriel and Dilek Hakkani-Tur

Thursday, June 6, 2019 (continued)

*How to Compare Summarizers without Target Length? Pitfalls, Solutions and Re-Examination of the Neural Summarization Literature*

Simeng Sun, Ori Shapira, Ido Dagan and Ani Nenkova

*BERT has a Mouth, and It Must Speak: BERT as a Markov Random Field Language Model*

Alex Wang and Kyunghyun Cho

*Neural Text Simplification in Low-Resource Conditions Using Weak Supervision*

Alessio Palmero Aprosio, Sara Tonelli, Marco Turchi, Matteo Negri and Mattia A. Di Gangi

*Paraphrase Generation for Semi-Supervised Learning in NLU*

Eunah Cho, He Xie and William M. Campbell

*Bilingual-GAN: A Step Towards Parallel Text Generation*

Ahmad Rashid, Alan Do Omri, Md Akmal Haidar, Qun Liu and Mehdi Reza-gholizadeh

*Learning Criteria and Evaluation Metrics for Textual Transfer between Non-Parallel Corpora*

Yuanzhe Pang and Kevin Gimpel

*Dual Supervised Learning for Natural Language Understanding and Generation*

Shang-Yu Su, Chao-Wei Huang and Yun-Nung Chen

*Designing a Symbolic Intermediate Representation for Neural Surface Realization*

Henry Elder, Jennifer Foster, James Barry and Alexander O'Connor

*Insertion-based Decoding with automatically Inferred Generation Order*

Jiatao Gu, Qi Liu and Kyunghyun Cho

*Neural Text Style Transfer via Denoising and Reranking*

Joseph Lee, Ziang Xie, Cindy Wang, Max Drach, Dan Jurafsky and Andrew Ng

*Generating Diverse Story Continuations with Controllable Semantics*

Lifu Tu, Xiaoan Ding, Dong Yu and Kevin Gimpel

*Better Automatic Evaluation of Open-Domain Dialogue Systems with Contextualized Embeddings*

Sarik Ghazarian, Johnny Wei, Aram Galstyan and Nanyun Peng

*Improved Zero-shot Neural Machine Translation via Ignoring Spurious Correlations*

Jiatao Gu, Yong Wang, Kyunghyun Cho and Victor O.K. Li

**Thursday, June 6, 2019 (continued)**

*Jointly Measuring Diversity and Quality in Text Generation Models*

Danial Alihosseini, Ehsan Montahaei and Mahdiah Soleymani Baghshah

**15:45–16:20** **Invited Talk: Tatsunori Hashimoto**

**16:20–16:55** **Invited Talk: Hal Daumé III**

**16:55–17:55** **Panel**

**17:55–18:00** *Closing Remarks*

