

ACL 2020

Neural Generation and Translation

Proceedings of the Fourth Workshop

July 5 - 10, 2020 Online (due to COVID-19 pandemic)

©2020 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

ISBN 978-1-952148-17-0

Introduction

Welcome to the Fourth Workshop on Neural Generation and Translation. This workshop aims to cultivate research on the leading edge in neural machine translation and other aspects of machine translation, generation, and multilinguality that utilize neural models. We received a total of 28 submissions in total. From the 21 long papers we accepted 16. There were two cross-submissions, three extended abstracts. All research papers were reviewed twice through a double blind review process, and avoiding conflicts of interest. The topics of the papers were split equally between the natural language generation and machine translation themes. We would like to thank all authors for their submissions, and the program committee members for their valuable efforts in reviewing the papers for the workshop.

Organizers:

Alexandra Birch, (Edinburgh)
Andrew Finch, (Apple)
Hiroaki Hayashi (CMU)
Kenneth Heafield (Edinburgh)
Marcin Junczys-Dowmunt (Microsoft)
Ioannis Konstas (Heriot Watt University)
Xian Li (Facebook)
Graham Neubig, (CMU)
Yusuke Oda, (Google)

Program Committee:

Roe Aharoni	Martin Andrews
Duygu Ataman	Joost Bastings
Alexandre Berard	Klinton Bicknell
Nikolay Bogoychev	Chris Brockett
Chris Brust	Daniel Cer
Boxing Chen	David Chiang
Nathan Dalal	Michael Denkowski
Li Dong	Kevin Duh
Ondřej Dušek	Marc Dymetman
Elozino Egonmwan	Andrew Finch
Markus Freitag	Claire Gardent
Kevin Gimpel	Isao Goto
Roman Grundkiewicz	Jiatao Gu
Felix Hieber	Hieu Hoang
Sébastien Jean	Mihir Kale
Yuta Kikuchi	Philipp Koehn
Ioannis Konstas	Shumpei Kubosawa
Sneha Kudugunta	Lemao Liu
Hongyin Luo	Benjamin Marie
Bill McDowell	Hideya Mino
Will Monroe	Makoto Morishita
Vivek Natarajan	Yusuke Oda
Vinay Rao	Andrew Runge
Chinnadhurai Sankar	Jean Senellart
Rico Sennrich	Alessandro Suglia
Simeng Sun	Alexey Tikonov
Yulia Tsvetkov	Ivan Vulić
Rui Wang	Xiaolin Wang
Taro Watanabe	Ruiyi Zhang
Biao Zhang	

Invited Speakers:

He He
Jiatao Gu
Shashi Narayan
Claire Gardent

Table of Contents

<i>Findings of the Fourth Workshop on Neural Generation and Translation</i> Kenneth Heafield, Hiroaki Hayashi, Yusuke Oda, Ioannis Konstas, Andrew Finch, Graham Neubig, Xian Li and Alexandra Birch	1
<i>Learning to Generate Multiple Style Transfer Outputs for an Input Sentence</i> Kevin Lin, Ming-Yu Liu, Ming-Ting Sun and Jan Kautz	10
<i>Balancing Cost and Benefit with Tied-Multi Transformers</i> Raj Dabre, Raphael Rubino and Atsushi Fujita	24
<i>Compressing Neural Machine Translation Models with 4-bit Precision</i> Alham Fikri Aji and Kenneth Heafield	35
<i>Meta-Learning for Few-Shot NMT Adaptation</i> Amr Sharaf, Hany Hassan and Hal Daumé III	43
<i>Automatically Ranked Russian Paraphrase Corpus for Text Generation</i> Vadim Gudkov, Olga Mitrofanova and Elizaveta Filippikh	54
<i>A Deep Reinforced Model for Zero-Shot Cross-Lingual Summarization with Bilingual Semantic Similarity Rewards</i> Zi-Yi Dou, Sachin Kumar and Yulia Tsvetkov	60
<i>A Question Type Driven and Copy Loss Enhanced Framework for Answer-Agnostic Neural Question Generation</i> Xiuyu Wu, Nan Jiang and Yunfang Wu	69
<i>A Generative Approach to Titling and Clustering Wikipedia Sections</i> Anjalie Field, Sascha Rothe, Simon Baumgartner, Cong Yu and Abe Ittycheriah	79
<i>The Unreasonable Volatility of Neural Machine Translation Models</i> Marzieh Fadaee and Christof Monz	88
<i>Leveraging Sentence Similarity in Natural Language Generation: Improving Beam Search using Range Voting</i> Sebastian Borgeaud and Guy Emerson	97
<i>Distill, Adapt, Distill: Training Small, In-Domain Models for Neural Machine Translation</i> Mitchell Gordon and Kevin Duh	110
<i>Training and Inference Methods for High-Coverage Neural Machine Translation</i> Michael Yang, Yixin Liu and Rahul Mayuranath	119
<i>Meeting the 2020 Duolingo Challenge on a Shoestring</i> Tadashi Nomoto	129
<i>English-to-Japanese Diverse Translation by Combining Forward and Backward Outputs</i> Masahiro Kaneko, Aizhan Imankulova, Toshio Hirasawa and Mamoru Komachi	134
<i>POSTECH Submission on Duolingo Shared Task</i> Junsu Park, Hongseok Kwon and Jong-Hyeok Lee	139

<i>The ADAPT System Description for the STAPLE 2020 English-to-Portuguese Translation Task</i> Rejwanul Haque, Yasmin Moslem and Andy Way	144
<i>Expand and Filter: CUNI and LMU Systems for the WNGT 2020 Duolingo Shared Task</i> Jindřich Libovický, Zdeněk Kasner, Jindřich Helcl and Ondřej Dušek	153
<i>Exploring Model Consensus to Generate Translation Paraphrases</i> Zhenhao Li, Marina Fomicheva and Lucia Specia	161
<i>Growing Together: Modeling Human Language Learning With n-Best Multi-Checkpoint Machine Translation</i> El Moatez Billah Nagoudi, Muhammad Abdul-Mageed and Hasan Cavusoglu	169
<i>Generating Diverse Translations via Weighted Fine-tuning and Hypotheses Filtering for the Duolingo STAPLE Task</i> Sweta Agrawal and Marine Carpuat	178
<i>The JHU Submission to the 2020 Duolingo Shared Task on Simultaneous Translation and Paraphrase for Language Education</i> Huda Khayrallah, Jacob Bremerman, Arya D. McCarthy, Kenton Murray, Winston Wu and Matt Post	188
<i>Simultaneous paraphrasing and translation by fine-tuning Transformer models</i> Rakesh Chada	198
<i>The NiuTrans System for WNGT 2020 Efficiency Task</i> Chi Hu, Bei Li, Yinqiao Li, Ye Lin, Yanyang Li, Chenglong Wang, Tong Xiao and Jingbo Zhu	204
<i>Efficient and High-Quality Neural Machine Translation with OpenNMT</i> Guillaume Klein, Dakun Zhang, Clément Chouteau, Josep Crego and Jean Senellart	211
<i>Edinburgh’s Submissions to the 2020 Machine Translation Efficiency Task</i> Nikolay Bogoychev, Roman Grundkiewicz, Alham Fikri Aji, Maximiliana Behnke, Kenneth Heafield, Sidharth Kashyap, Emmanouil-Ioannis Farsarakis and Mateusz Chudyk	218
<i>Improving Document-Level Neural Machine Translation with Domain Adaptation</i> Sami Ul Haq, Sadaf Abdul Rauf, Arslan Shoukat and Noor-e- Hira	225
<i>Simultaneous Translation and Paraphrase for Language Education</i> Stephen Mayhew, Klinton Bicknell, Chris Brust, Bill McDowell, Will Monroe and Burr Settles	232

Workshop Program

Findings of the Fourth Workshop on Neural Generation and Translation

Kenneth Heafield, Hiroaki Hayashi, Yusuke Oda, Ioannis Konstas, Andrew Finch, Graham Neubig, Xian Li and Alexandra Birch

Learning to Generate Multiple Style Transfer Outputs for an Input Sentence

Kevin Lin, Ming-Yu Liu, Ming-Ting Sun and Jan Kautz

Balancing Cost and Benefit with Tied-Multi Transformers

Raj Dabre, Raphael Rubino and Atsushi Fujita

Compressing Neural Machine Translation Models with 4-bit Precision

Alham Fikri Aji and Kenneth Heafield

Meta-Learning for Few-Shot NMT Adaptation

Amr Sharaf, Hany Hassan and Hal Daumé III

Automatically Ranked Russian Paraphrase Corpus for Text Generation

Vadim Gudkov, Olga Mitrofanova and Elizaveta Filippskikh

Increasing Lexical Diversity in Plug and Play Language Models

Soham Parikh, Daphne Ippolito and Satyarth Vaidya

A Deep Reinforced Model for Zero-Shot Cross-Lingual Summarization with Bilingual Semantic Similarity Rewards

Zi-Yi Dou, Sachin Kumar and Yulia Tsvetkov

A Question Type Driven and Copy Loss Enhanced Framework for Answer-Agnostic Neural Question Generation

Xiuyu Wu, Nan Jiang and Yunfang Wu

When and Why is Unsupervised Neural Machine Translation Useless?

Yunsu Kim, Miguel Graça and Hermann Ney

A Generative Approach to Titling and Clustering Wikipedia Sections

Anjalie Field, Sascha Rothe, Simon Baumgartner, Cong Yu and Abe Ittycheriah

The Unreasonable Volatility of Neural Machine Translation Models

Marzieh Fadaee and Christof Monz

No Day Set (continued)

Leveraging Sentence Similarity in Natural Language Generation: Improving Beam Search using Range Voting

Sebastian Borgeaud and Guy Emerson

Transformers without Tears: Improving the Normalization of Self-Attention

Toan Q. Nguyen and Julian Salazar

Masked Language Model Scoring

Julian Salazar, Davis Liang, Toan Q. Nguyen and Katrin Kirchhoff

Distill, Adapt, Distill: Training Small, In-Domain Models for Neural Machine Translation

Mitchell Gordon and Kevin Duh

Improving Neural Machine Translation Using Energy-Based Models

Subhajit Naskar, Amirmohammad Rooshenas and Andrew McCallum

Training and Inference Methods for High-Coverage Neural Machine Translation

Michael Yang, Yixin Liu and Rahul Mayuranath

Meeting the 2020 Duolingo Challenge on a Shoestring

Tadashi Nomoto

English-to-Japanese Diverse Translation by Combining Forward and Backward Outputs

Masahiro Kaneko, Aizhan Imankulova, Tosho Hirasawa and Mamoru Komachi

POSTECH Submission on Duolingo Shared Task

Junsu Park, Hongseok Kwon and Jong-Hyeok Lee

The ADAPT System Description for the STAPLE 2020 English-to-Portuguese Translation Task

Rejwanul Haque, Yasmin Moslem and Andy Way

Expand and Filter: CUNI and LMU Systems for the WNGT 2020 Duolingo Shared Task

Jindřich Libovický, Zdeněk Kasner, Jindřich Helcl and Ondřej Dušek

Exploring Model Consensus to Generate Translation Paraphrases

Zhenhao Li, Marina Fomicheva and Lucia Specia

No Day Set (continued)

Growing Together: Modeling Human Language Learning With n-Best Multi-Checkpoint Machine Translation

El Moatez Billah Nagoudi, Muhammad Abdul-Mageed and Hasan Cavusoglu

Generating Diverse Translations via Weighted Fine-tuning and Hypotheses Filtering for the Duolingo STAPLE Task

Sweta Agrawal and Marine Carpuat

The JHU Submission to the 2020 Duolingo Shared Task on Simultaneous Translation and Paraphrase for Language Education

Huda Khayrallah, Jacob Bremerman, Arya D. McCarthy, Kenton Murray, Winston Wu and Matt Post

Simultaneous paraphrasing and translation by fine-tuning Transformer models

Rakesh Chada

The NiuTrans System for WNGT 2020 Efficiency Task

Chi Hu, Bei Li, Yinqiao Li, Ye Lin, Yanyang Li, Chenglong Wang, Tong Xiao and Jingbo Zhu

Efficient and High-Quality Neural Machine Translation with OpenNMT

Guillaume Klein, Dakun Zhang, Clément Chouteau, Josep Crego and Jean Senellart

Edinburgh's Submissions to the 2020 Machine Translation Efficiency Task

Nikolay Bogoychev, Roman Grundkiewicz, Alham Fikri Aji, Maximiliana Behnke, Kenneth Heafield, Sidharth Kashyap, Emmanouil-Ioannis Farsarakis and Mateusz Chudyk

Improving Document-Level Neural Machine Translation with Domain Adaptation

Sami Ul Haq, Sadaf Abdul Rauf, Arslan Shoukat and Noor-e- Hira

Simultaneous Translation and Paraphrase for Language Education

Stephen Mayhew, Klinton Bicknell, Chris Brust, Bill McDowell, Will Monroe and Burr Settles

