

ACL 2018

Neural Machine Translation and Generation

Proceedings of the Second Workshop

July 20, 2018
Melbourne, Australia



©2018 The Association for Computational Linguistics

Apple and the Apple logo are trademarks of Apple Inc., registered in the U.S. and other countries.

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 978-1-948087-40-7

Introduction

Machine Translation and Generation focusing on Neural Machine Translation (NMT) technology. 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. In this year's workshop we are extremely pleased to be able to host four invited talks from leading lights in the field, namely: Jacob Devlin, Rico Sennrich, Jason Weston, and Yulia Tsvetkov. In addition this year's workshop will feature a session devoted to a new shared task on efficient machine translation. We received a total of 25 submissions, and accepted 16 for inclusion in the workshop, an acceptance rate of 64%. Due to the large number of invited talks, and to encourage discussion, only the two papers selected for best paper awards will be presented orally, and the remainder will be presented in a single poster session. 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. We would also like to thank Amazon, Apple and Google for their generous sponsorship.

Organizers:

Alexandra Birch, (Edinburgh)
Andrew Finch, (Apple)
Thang Luong, (Google)
Graham Neubig, (CMU)
Yusuke Oda, (Google)

Program Committee:

Roe Aharoni, (Bar Ilan University)
Joost Bastings, (University of Amsterdam)
Yonatan Belinkov, (Qatar Computing Research Institute)
Marine Carpuat, (University of Maryland)
Boxing Chen, (National Research Council, Canada)
Eunah Cho, (Amazon)
Michael Denkowski, (Amazon)
Kevin Duh, (JHU)
Cong Duy Vu Hoang, (University of Melbourne)
Markus Freitag, (Google)
Isao Goto, (NHK)
Jiatao Gu, (The University of Hong Kong)
Barry Haddow, (Edinburgh)
Sebastien Jean, (Montreal)
Yuta Kikuchi, (Preferred Networks)
Philipp Koehn, (Johns Hopkins University)
Yannis Konstas, (Heriot-Watt University)
Shumpei Kubosawa, (NEC)
Mirella Lapata, (University of Edinburgh)
Shujie Liu, (Microsoft)
Lemao Liu, (Tencent AI Lab)
Haitao Mi, (Ant Financial US)
Hideya Mino, (NHK)
Makoto Morishita, (Nara Institute of Science and Technology)
Preslav Nakov, (QCRI)
Hieu Pham, (Google)
Alexander Rush, (Harvard)
Abigail See, (Stanford)
Rico Sennrich, (Edinburgh)
Raphael Shu, (The University of Tokyo)
Akihiro Tamura, (Ehime University)
Rui Wang, (NICT)
Xiaolin Wang, (NICT)
Taro Watanabe, (Google)
Xingxing Zhang, (University of Edinburgh)

Invited Speakers:

Jacob Devlin, (Google)
Rico Sennrich, (Edinburgh)
Yulia Tsvetkov, (CMU)
Jason Weston, (Facebook)

Table of Contents

<i>Findings of the Second Workshop on Neural Machine Translation and Generation</i>	
Alexandra Birch, Andrew Finch, Minh-Thang Luong, Graham Neubig and Yusuke Oda	1
<i>A Shared Attention Mechanism for Interpretation of Neural Automatic Post-Editing Systems</i>	
Inigo Jauregi Unanue, Ehsan Zare Borzeshi and Massimo Piccardi	11
<i>Iterative Back-Translation for Neural Machine Translation</i>	
Vu Cong Duy Hoang, Philipp Koehn, Gholamreza Haffari and Trevor Cohn	18
<i>Inducing Grammars with and for Neural Machine Translation</i>	
Yonatan Bisk and Ke Tran	25
<i>Regularized Training Objective for Continued Training for Domain Adaptation in Neural Machine Translation</i>	
Huda Khayrallah, Brian Thompson, Kevin Duh and Philipp Koehn	36
<i>Controllable Abstractive Summarization</i>	
Angela Fan, David Grangier and Michael Auli	45
<i>Enhancement of Encoder and Attention Using Target Monolingual Corpora in Neural Machine Translation</i>	
Kenji Imamura, Atsushi Fujita and Eiichiro Sumita	55
<i>Document-Level Adaptation for Neural Machine Translation</i>	
Sachith Sri Ram Kothur, Rebecca Knowles and Philipp Koehn	64
<i>On the Impact of Various Types of Noise on Neural Machine Translation</i>	
Huda Khayrallah and Philipp Koehn	74
<i>Bi-Directional Neural Machine Translation with Synthetic Parallel Data</i>	
Xing Niu, Michael Denkowski and Marine Carpuat	84
<i>Multi-Source Neural Machine Translation with Missing Data</i>	
Yuta Nishimura, Katsuhito Sudoh, Graham Neubig and Satoshi Nakamura	92
<i>Towards one-shot learning for rare-word translation with external experts</i>	
Ngoc-Quan Pham, Jan Niehues and Alexander Waibel	100
<i>NICT Self-Training Approach to Neural Machine Translation at NMT-2018</i>	
Kenji Imamura and Eiichiro Sumita	110
<i>Fast Neural Machine Translation Implementation</i>	
Hieu Hoang, Tomasz Dwojak, Rihards Krislauks, Daniel Torregrosa and Kenneth Heafield	116
<i>OpenNMT System Description for WNMT 2018: 800 words/sec on a single-core CPU</i>	
Jean Senellart, Dakun Zhang, Bo WANG, Guillaume KLEIN, Jean-Pierre Ramatchandirin, Josep Crego and Alexander Rush	122
<i>Marian: Cost-effective High-Quality Neural Machine Translation in C++</i>	
Marcin Junczys-Dowmunt, Kenneth Heafield, Hieu Hoang, Roman Grundkiewicz and Anthony Aue	129

Workshop Program

July 20, 2018

09:00–09:10 *Welcome and Opening Remarks*

Findings of the Second Workshop on Neural Machine Translation and Generation
Alexandra Birch, Andrew Finch, Minh-Thang Luong, Graham Neubig and Yusuke Oda

09:10–10:00 *Keynote 1*
Jacob Devlin

10:00–10:30 *Shared Task Overview*

10:30–11:00 *Coffee Break*

11:00–11:30 *Marian: Fast Neural Machine Translation in C++*

11:30–12:20 *Keynote 2*
Rico Sennrich

12:20–13:20 *Lunch Break*

13:20–13:50 **Best Paper Session**

13:50–14:40 *Keynote 3*
Jason Weston

14:40–15:30 *Keynote 4*
Yulia Tsvetkov

15:30–16:00 *Coffee Break*

July 20, 2018 (continued)

16:00–17:30 Poster Session

A Shared Attention Mechanism for Interpretation of Neural Automatic Post-Editing Systems

Inigo Jauregi Unanue, Ehsan Zare Borzeshi and Massimo Piccardi

Iterative Back-Translation for Neural Machine Translation

Vu Cong Duy Hoang, Philipp Koehn, Gholamreza Haffari and Trevor Cohn

Inducing Grammars with and for Neural Machine Translation

Yonatan Bisk and Ke Tran

Regularized Training Objective for Continued Training for Domain Adaptation in Neural Machine Translation

Huda Khayrallah, Brian Thompson, Kevin Duh and Philipp Koehn

Controllable Abstractive Summarization

Angela Fan, David Grangier and Michael Auli

Enhancement of Encoder and Attention Using Target Monolingual Corpora in Neural Machine Translation

Kenji Imamura, Atsushi Fujita and Eiichiro Sumita

Document-Level Adaptation for Neural Machine Translation

Sachith Sri Ram Kothur, Rebecca Knowles and Philipp Koehn

On the Impact of Various Types of Noise on Neural Machine Translation

Huda Khayrallah and Philipp Koehn

Bi-Directional Neural Machine Translation with Synthetic Parallel Data

Xing Niu, Michael Denkowski and Marine Carpuat

Multi-Source Neural Machine Translation with Missing Data

Yuta Nishimura, Katsuhito Sudoh, Graham Neubig and Satoshi Nakamura

Towards one-shot learning for rare-word translation with external experts

Ngoc-Quan Pham, Jan Niehues and Alexander Waibel

July 20, 2018 (continued)

NICT Self-Training Approach to Neural Machine Translation at NMT-2018

Kenji Imamura and Eiichiro Sumita

Fast Neural Machine Translation Implementation

Hieu Hoang, Tomasz Dwojak, Rihards Krislauks, Daniel Torregrosa and Kenneth Heafield

OpenNMT System Description for WNMT 2018: 800 words/sec on a single-core CPU

Jean Senellart, Dakun Zhang, Bo WANG, Guillaume KLEIN, Jean-Pierre Ramatchandirin, Josep Crego and Alexander Rush

Marian: Cost-effective High-Quality Neural Machine Translation in C++

Marcin Junczys-Dowmunt, Kenneth Heafield, Hieu Hoang, Roman Grundkiewicz and Anthony Aue

On Individual Neurons in Neural Machine Translation

D. Anthony Bau, Yonatan Belinkov, Hassan Sajjad, Nadir Durrani, Fahim Dalvi and James Glass

Parameter Sharing Strategies in Neural Machine Translation

Sébastien Jean, Stanislas Lauly and Kyunghyun Cho

Modeling Latent Sentence Structure in Neural Machine Translation

Joost Bastings, Wilker Aziz, Ivan Titov and Khalil Simaan

Extreme Adaptation for Personalized Neural Machine Translation

Paul Michel and Graham Neubig

Exploiting Semantics in Neural Machine Translation with Graph Convolutional Networks

Diego Marcheggiani, Joost Bastings and Ivan Titov

17:30–17:40 *Closing Remarks*

