IWSLT 2020

The 17th International Conference on Spoken Language Translation

Proceedings of the Conference

July 9 - 10, 2020

©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 978-1-952148-07-1

Preface

The International Conference on Spoken Language Translation (IWSLT) is the premiere annual scientific conference for the study, development and evaluation of spoken language translation technology. Launched in 2004 and spun out from the C-STAR speech translation consortium before it (1992-2003), IWSLT is the main venue for scientific exchange on all topics related to speech-to-text translation, speech-to- speech translation, simultaneous and consecutive translation, speech dubbing, cross-lingual communication including all multimodal, emotional, paralinguistic, and stylistic aspects and their applications in the field. The conference organizes evaluations and workshop sessions around challenge areas, and presents scientific work and system descriptions.

In 2020, IWSLT joins ACL as the first ACL speech translation conference. It features six challenge tracks: 1.) Simultaneous speech translation, 2.) Video speech translation, 3.) Offline speech translation, 4.) Conversational speech translation, 5.) Open domain translation, and 6.) Non-native speech translation. These topics represent open problems toward effective cross-lingual communication and we expect the community effort and discussion will greatly advance the state of the field. Each track was coordinated by a chair. The resulting evaluation campaigns attracted a total of 30 teams, from academy and industry. System submissions resulted in system papers that will be presented at the conference.

The great turnout this year demonstrates the growing academic and commercial interest in spoken language translation. It has been empowered by growing availability of data resources also facilitated by the conference over the years. The versatility of neural network models and broad availability of training frameworks further support this trend as even smaller groups can now tackle complex problems, like speech translation, directly. Given the growing interest and international demand, we invite researchers and potential organizers to propose new IWSLT challenges for the future.

Following our call for papers this year, 43 submissions were received. In a blind review process, 9 research papers were selected out of 19 for oral presentation (47%) in addition to 24 system papers. The program committee is excited about the quality of the accepted papers and expects lively discussion and exchange at the conference.

The conference chairs and organizers would like to express their gratitude to everyone who contributed and supported IWSLT. Our IWSLT-20 program exceeds all our expectations in quality and breath, particularly when considering the challenges during a pandemic under lock-downs and health and travel restrictions. We thank the challenge track chairs, organizers, and participants, the program chairs and committee members, as well as all the authors that went the extra mile to submit system and research papers to IWSLT, and make this year's conference our most vibrant than ever. We also wish to express our sincere gratitude to ACL for hosting our conference and for arranging the logistics and infrastructure that allow us to hold IWSLT 2020 as a virtual online conference.

Welcome to IWSLT 2020, in Seattle, Mountain View or wherever you may be hiding out!

Marcello Federico and Alex Waibel
Conference Chairs

Organizers:

Marcello Federico, Amazon, USA (Chair)
Alex Waibel, CMU, USA (Chair)
Kevin Knight, DiDi Labs, USA
Satoshi Nakamura, NAIST, Japan
Hermann Ney, RWTH, Germany
Jan Niehues, Maastricht University, Netherlands
Sebastian Stüker, KIT, Germany
Dekai Wu, HKUST Hong Kong, China
Joseph Mariani, LIMSI/CNRS, France
Francois Yvon, LIMSI/CNRS, France

Challenge Track Organizers:

Simultaneous Speech Translation
Jiatao Gu, Facebook, USA (Chair)
Juan Pino, Facebook, USA
Changhan Wang, Facebook, USA
Xutai Ma, JHU, USA
Fahim Dalvi, QCRI, Qatar
Nadir Durrani, QCRI, Qatar
Video Speech Translation
Fei Huang, Alibaba, USA (Chair)
Nguyen Bach, Alibaba, USA
Wei Luo, Alibaba, USA
Boxing Chen, Alibaba, USA
Fei Huang, Alibaba, USA
Offline Speech Translation
Marco Turchi, FBK, Italy (Chair)
Matteo Negri, FBK, Italy
Roldano Cattoni, FBK, Italy
Sebastian Stüker, KIT, Germany
Jan Niehues, Maastricht University, Netherland
Conversational Speech Translation
Elizabeth Salesky, JHU, USA (Chair)
Open Domain Translation
Ajay Nagesh, Didi Labs, USA (Chair)
Amittai Axelrod, Didi Labs, USA
Arkady Arkhangorodsky, Didi Labs, USA
Boliang Zhang, Didi Labs, USA
Xing Shi, Didi Labs, USA
Non-Native Speech Translation
Ondrej Bojar, Charles University, Czech Republic (Chair)
Ebrahim Ansari, Charles University, Czech Republic

Program Committee:

Christian Federmann, Microsoft Research, USA (Chair)
Will Lewis, Microsoft Research, USA (Chair)
Akiko Eriguchi, Microsoft, USA
Antonio Toral, U. Groningen, Netherlands
Boxing Chen, Alibaba, China
Christian Hardmeier, Uppsala U., Sweden
David Vilar, Amazon, Germany
Dongdong Zhang, Microsoft, USA
Duygu Ataman, U. Zurich, Switzerland
Elizabeth Salesky, JHU, USA
Evgeny Matusov, AppTek, Germany
Frank Seide, Microsoft, USA
Hieu Hoang, U. Edinburgh, USA
Hoang Cuong, Amazon, USA
José G. C. de Souza, eBay, Germany
Juan Pino, Facebook, USA
Julian Salazar, Amazon, USA
Katsuhito Sudoh, NAIST, Japan
Keisuke Sakaguchi, AI2 Seattle, USA
Krzysztof Wolk, PJA, Poland
Laurent Besacier, IMAG, France
Marcin Junczys-Dowmunt, Microsoft, USA
Marco Turchi, FBK, Italy
Markus Freitag, Google, USA
Matteo Negri, FBK, Italy
Matthias Huck, LMU, Germany
Mattia di Gangi, FBK/U. Trento, Italy
Nguyen Bach, Alibaba, USA
Nicholas Ruiz, Interaction, USA
Preslav Nakov, QCRI, Qatar
Roman Grundkiewicz, U. Edinburgh, UK
Surafel Melaku Lakew, FBK/U. Trento, Italy
Thanh-Le Ha, KIT, Germany
Vincent Vandeghinste, KU Leuven, Belgium
Young Jin Kim, Microsoft, USA
Yves Lepage, U. Waseda, Japan

Table of Contents

FINDINGS OF THE IWSLT 2020 EVALUATION CAMPAIGN
Ebrahim Ansari, amittai axelrod, Nguyen Bach, Ondřej Bojar, Roldano Cattoni, Fahim Dalvi, Nadir Durrani, Marcello Federico, Christian Federmann, Jiatao Gu, Fei Huang, Kevin Knight, Xutai Ma, Ajay Nagesh, Matteo Negri, Jan Niehues, Juan Pino, Elizabeth Salesky, Xing Shi, Sebastian Stüker, Marco Turchi, Alexander Waibel and Changhan Wang 1

ON-TRAC Consortium for End-to-End and Simultaneous Speech Translation Challenge Tasks at IWSLT 2020
Maha Elbayad, Ha Nguyen, Fethi Bougares, Natalia Tomashenko, Antoine Caubrière, Benjamin Lecouteux, Yannick Estève and Laurent Besacier 35

Start-Before-End and End-to-End: Neural Speech Translation by AppTek and RWTH Aachen University
Parnia Bahar, Patrick Wilken, Tamer Alkhouli, Andreas Guta, Pavel Golik, Evgeny Matusov and Christian Herold 44

KIT's IWSLT 2020 SLT Translation System
Ngoc-Quan Pham, Felix Schneider, Tuan-Nam Nguyen, Thanh-Le Ha, Thai Son Nguyen, Maximilian Awiszus, Sebastian Stüker and Alexander Waibel 55

End-to-End Simultaneous Translation System for IWSLT2020 Using Modality Agnostic Meta-Learning
Hou Jeung Han, Mohd Abbas Zaidi, Sathish Reddy Indurthi, Nikhil Kumar Lakumarapu, Beomseok Lee and Sangha Kim 62

DiDi Labs' End-to-end System for the IWSLT 2020 Offline Speech TranslationTask
Arkady Arkhangorodsky, Yiqi Huang and amittai axelrod 69

End-to-End Offline Speech Translation System for IWSLT 2020 using Modality Agnostic Meta-Learning
Nikhil Kumar Lakumarapu, Beomseok Lee, Sathish Reddy Indurthi, Hou Jeung Han, Mohd Abbas Zaidi and Sangha Kim 73

End-to-End Speech-Translation with Knowledge Distillation: FBK@IWSLT2020
Marco Gaido, Mattia A. Di Gangi, Matteo Negri and Marco Turchi 80

SRPOL's System for the IWSLT 2020 End-to-End Speech Translation Task
Tomasz Potapczyk and Pawel Przybysz 89

The University of Helsinki Submission to the IWSLT2020 Offline SpeechTranslation Task
Raúl Vázquez, Mikko Aulamo, Umut Sulubacak and Jörg Tiedemann 95

The AFRL IWSLT 2020 Systems: Work-From-Home Edition
Brian Ore, Eric Hansen, Tim Anderson and Jeremy Gwinnup 103

LIT Team's System Description for Japanese-Chinese Machine Translation Task in IWSLT 2020
Yimeng Zhuang, Yuan Zhang and Lijie Wang 109

OPPO's Machine Translation System for the IWSLT 2020 Open Domain Translation Task
Qian Zhang, Xiaopu Li, Dawei Dang, Tingxun Shi, Di Ai, Zhengshan Xue and Jie Hao 114

Character Mapping and Ad-hoc Adaptation: Edinburgh's IWSLT 2020 Open Domain Translation System
Pinzhen Chen, Nikolay Bogoychev and Ulrich Germann 122

CASIA's System for IWSLT 2020 Open Domain Translation
Qian Wang, Yuchen Liu, Cong Ma, Yu Lu, Yining Wang, Long Zhou, Yang Zhao, Jiajun Zhang and Chengqing Zong 130

Deep Blue Sonics' Submission to IWSLT 2020 Open Domain Translation Task
Enmin Su and Yi Ren 140

University of Tsukuba's Machine Translation System for IWSLT20 Open Domain Translation Task
Hongyi Cui, Yizhen Wei, Shohei Iida, Takehito Utsuro and Masaaki Nagata 145

Xiaomi's Submissions for IWSLT 2020 Open Domain Translation Task
Yuhui Sun, Mengxue Guo, Xiang Li, Jianwei Cui and Bin Wang 149

ISTIC's Neural Machine Translation System for IWSLT'2020
jiaze wei, wenbin liu, zhenfeng wu, you pan and yanqing he 158

Octanove Labs' Japanese-Chinese Open Domain Translation System
Masato Hagiwara 166

NAIST's Machine Translation Systems for IWSLT 2020 Conversational Speech Translation Task
Ryo Fukuda, Katsuhito Sudoh and Satoshi Nakamura 172

Generating Fluent Translations from Disfluent Text Without Access to Fluent References: IIT Bombay@IWSLT2020
Nikhil Saini, Jyotsana Khatri, Preethi Jyothi and Pushpak Bhattacharyya 178

The HW-TSC Video Speech Translation System at IWSLT 2020
Minghan Wang, Hao Yang, Yao Deng, Ying Qin, Lizhi Lei, Daimeng Wei, Hengchao Shang, Ning Xie, Xiaochun Li and Jiaxian Guo 187

CUNI Neural ASR with Phoneme-Level Intermediate Step for Non-Native SLT at IWSLT 2020
Peter Polák, Sangeet Sagar, Dominik Macháček and Ondřej Bojar 191

ELITR Non-Native Speech Translation at IWSLT 2020
Dominik Macháček, Jonáš Kratochvíl, Sangeet Sagar, Matúš Žilinec, Ondřej Bojar, Thai-Son Nguyen, Felix Schneider, Philip Williams and Yuekun Yao 200

Is 42 the Answer to Everything in Subtitling-oriented Speech Translation?
Alina Karakanta, Matteo Negri and Marco Turchi 209

Re-translation versus Streaming for Simultaneous Translation
Naveen Arivazhagan, Colin Cherry, Wolfgang Macherey and George Foster 220

Towards Stream Translation: Adaptive Computation Time for Simultaneous Machine Translation
Felix Schneider and Alexander Waibel 228

Neural Simultaneous Speech Translation Using Alignment-Based Chunking
Patrick Wilken, Tamer Alkhouli, Evgeny Matusov and Pavel Golik 237

Adapting End-to-End Speech Recognition for Readable Subtitles
Danni Liu, Jan Niehues and Gerasimos Spanakis 247

From Speech-to-Speech Translation to Automatic Dubbing
Marcello Federico, Robert Enyedi, Roberto Barra-Chicote, Ritwik Giri, Umut Isik, Arvindh Krishnaswamy and Hassan Sawaf 257

Joint Translation and Unit Conversion for End-to-end Localization
Georgiana Dinu, Prashant Mathur, Marcello Federico, Stanislas Lauly and Yaser Al-Onaizan . . 265

Efficient Automatic Punctuation Restoration Using Bidirectional Transformers with Robust Inference
Maury Courtland, Adam Faulkner and Gayle McElvain 272

How Human is Machine Translationese? Comparing Human and Machine Translations of Text and Speech
Yuri Bizzoni, Tom S Juzek, Cristina España-Bonet, Koel Dutta Chowdhury, Josef van Genabith and Elke Teich 280

Conference Program

Day 1

Evaluation Overview

FINDINGS OF THE IWSLT 2020 EVALUATION CAMPAIGN
Ebrahim Ansari, amittai axelrod, Nguyen Bach, Ondřej Bojar, Roldano Cattoni, Fahim Dalvi, Nadir Durrani, Marcello Federico, Christian Federmann, Jiatao Gu, Fei Huang, Kevin Knight, Xutai Ma, Ajay Nagesh, Matteo Negri, Jan Niehues, Juan Pino, Elizabeth Salesky, Xing Shi, Sebastian Stüker, Marco Turchi, Alexander Waibel and Changhan Wang

Day 1-2

System Papers: Simultaneous Speech Translation

ON-TRAC Consortium for End-to-End and Simultaneous Speech Translation Challenge Tasks at IWSLT 2020
Maha Elbayad, Ha Nguyen, Fethi Bougares, Natalia Tomashenko, Antoine Caubrière, Benjamin Lecouteux, Yannick Estève and Laurent Besacier

Start-Before-End and End-to-End: Neural Speech Translation by AppTek and RWTH Aachen University
Parnia Bahar, Patrick Wilken, Tamer Alkhouli, Andreas Guta, Pavel Golik, Evgeny Matusov and Christian Herold

KIT's IWSLT 2020 SLT Translation System
Ngoc-Quan Pham, Felix Schneider, Tuan-Nam Nguyen, Thanh-Le Ha, Thai Son Nguyen, Maximilian Awiszus, Sebastian Stüker and Alexander Waibel

End-to-End Simultaneous Translation System for IWSLT2020 Using Modality Agnostic Meta-Learning
Hou Jeung Han, Mohd Abbas Zaidi, Sathish Reddy Indurthi, Nikhil Kumar Lakumarapu, Beomseok Lee and Sangha Kim

Day 1-2 (continued)

System Papers: Offline Speech Translation

DiDi Labs' End-to-end System for the IWSLT 2020 Offline Speech TranslationTask
Arkady Arkhangorodsky, Yiqi Huang and amittai axelrod

End-to-End Offline Speech Translation System for IWSLT 2020 using Modality Agnostic Meta-Learning
Nikhil Kumar Lakumarapu, Beomseok Lee, Sathish Reddy Indurthi, Hou Jeung Han, Mohd Abbas Zaidi and Sangha Kim

End-to-End Speech-Translation with Knowledge Distillation: FBK@IWSLT2020
Marco Gaido, Mattia A. Di Gangi, Matteo Negri and Marco Turchi

SRPOL's System for the IWSLT 2020 End-to-End Speech Translation Task
Tomasz Potapczyk and Pawel Przybysz

The University of Helsinki Submission to the IWSLT2020 Offline SpeechTranslation Task
Raúl Vázquez, Mikko Aulamo, Umut Sulubacak and Jörg Tiedemann

The AFRL IWSLT 2020 Systems: Work-From-Home Edition
Brian Ore, Eric Hansen, Tim Anderson and Jeremy Gwinnup

System Papers: Open Domain Translation

LIT Team's System Description for Japanese-Chinese Machine Translation Task in IWSLT 2020
Yimeng Zhuang, Yuan Zhang and Lijie Wang

OPPO's Machine Translation System for the IWSLT 2020 Open Domain Translation Task
Qian Zhang, Xiaopu Li, Dawei Dang, Tingxun Shi, Di Ai, Zhengshan Xue and Jie Hao

Character Mapping and Ad-hoc Adaptation: Edinburgh's IWSLT 2020 Open Domain Translation System
Pinzhen Chen, Nikolay Bogoychev and Ulrich Germann

CASIA's System for IWSLT 2020 Open Domain Translation
Qian Wang, Yuchen Liu, Cong Ma, Yu Lu, Yining Wang, Long Zhou, Yang Zhao, Jiajun Zhang and Chengqing Zong

Day 1-2 (continued)

Deep Blue Sonics' Submission to IWSLT 2020 Open Domain Translation Task
Enmin Su and Yi Ren

University of Tsukuba's Machine Translation System for IWSLT20 Open Domain Translation Task
Hongyi Cui, Yizhen Wei, Shohei Iida, Takehito Utsuro and Masaaki Nagata

Xiaomi's Submissions for IWSLT 2020 Open Domain Translation Task
Yuhui Sun, Mengxue Guo, Xiang Li, Jianwei Cui and Bin Wang

ISTIC's Neural Machine Translation System for IWSLT'2020
jiaze wei, wenbin liu, zhenfeng wu, you pan and yanqing he

Octanove Labs' Japanese-Chinese Open Domain Translation System
Masato Hagiwara

System Papers: Conversational Speech Translation

NAIST's Machine Translation Systems for IWSLT 2020 Conversational Speech Translation Task
Ryo Fukuda, Katsuhito Sudoh and Satoshi Nakamura

Generating Fluent Translations from Disfluent Text Without Access to Fluent References: IIT Bombay@IWSLT2020
Nikhil Saini, Jyotsana Khatri, Preethi Jyothi and Pushpak Bhattacharyya

Day 1-2 (continued)

System Papers: Video Speech Translation

The HW-TSC Video Speech Translation System at IWSLT 2020
Minghan Wang, Hao Yang, Yao Deng, Ying Qin, Lizhi Lei, Daimeng Wei, Hengchao Shang, Ning Xie, Xiaochun Li and Jiaxian Guo

System Papers: Non-Native Speech Translation

CUNI Neural ASR with Phoneme-Level Intermediate Step for Non-Native SLT at IWSLT 2020
Peter Polák, Sangeet Sagar, Dominik Macháček and Ondřej Bojar

ELITR Non-Native Speech Translation at IWSLT 2020
Dominik Macháček, Jonáš Kratochvíl, Sangeet Sagar, Matúš Žilinec, Ondřej Bojar, Thai-Son Nguyen, Felix Schneider, Philip Williams and Yuekun Yao

Research Papers

Is 42 the Answer to Everything in Subtitling-oriented Speech Translation?
Alina Karakanta, Matteo Negri and Marco Turchi

Re-translation versus Streaming for Simultaneous Translation
Naveen Arivazhagan, Colin Cherry, Wolfgang Macherey and George Foster

Towards Stream Translation: Adaptive Computation Time for Simultaneous Machine Translation
Felix Schneider and Alexander Waibel

Neural Simultaneous Speech Translation Using Alignment-Based Chunking
Patrick Wilken, Tamer Alkhouli, Evgeny Matusov and Pavel Golik

Adapting End-to-End Speech Recognition for Readable Subtitles
Danni Liu, Jan Niehues and Gerasimos Spanakis

From Speech-to-Speech Translation to Automatic Dubbing
Marcello Federico, Robert Enyedi, Roberto Barra-Chicote, Ritwik Giri, Umut Isik, Arvindh Krishnaswamy and Hassan Sawaf

Day 1-2 (continued)

Joint Translation and Unit Conversion for End-to-end Localization
Georgiana Dinu, Prashant Mathur, Marcello Federico, Stanislas Lauly and Yaser Al-Onaizan

Efficient Automatic Punctuation Restoration Using Bidirectional Transformers with Robust Inference
Maury Courtland, Adam Faulkner and Gayle McElvain

How Human is Machine Translationese? Comparing Human and Machine Translations of Text and Speech
Yuri Bizzoni, Tom S Juzek, Cristina España-Bonet, Koel Dutta Chowdhury, Josef van Genabith and Elke Teich