

DSNNLG 2019

1st Workshop on Discourse Structure in Neural NLG

Proceedings of the Workshop

November 1, 2019
Tokyo, Japan

Endorsed by SIGGEN and SIGDIAL

Sponsored by Saarland University (SFB1102 and SFB-TR 248) and The Ohio State University

©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

ISBN 978-1-950737-67-3

Introduction

Welcome to the 1st Workshop on Discourse Structure in Neural NLG, a workshop held in conjunction with INLG 2019, the International Conference on Natural Language Generation, in Tokyo, Japan.

Neural methods for natural language generation (NNLG) arrived with much fanfare a few years ago and became the dominant method employed in the E2E NLG Challenge. While neural methods promise flexible, end-to-end trainable models, recent studies have revealed their inability to produce satisfactory output for longer or more complex texts as well as how the black-box nature of these models makes them difficult to control, in contrast to traditional NLG architectures that make use of explicit representations of discourse structure and/or sentence planning operations. As such, several papers have recently appeared that explore how to incorporate intermediate structures into NNLG or otherwise improve coherence and cohesion.

This workshop aims to encourage further research on enhancing quality in NNLG in terms of discourse coherence and cohesion along with ways to make NNLG models easier to control. Topics covered will include the limits of current end-to-end NNLG with respect to sentence planning and discourse structure; methods for improving discourse coherence and cohesion in NNLG, for example by making better use of discourse connectives, or by avoiding unnecessary repetition; methods for control and interpretability of NNLG, for example by providing more explicit guidance or structure in the input; and better methods for evaluating discourse coherence and cohesion in NNLG.

These proceedings include a total of four papers, chosen from seven submitted papers, each reviewed by three members of the program committee. In addition to presentation of papers, the workshop will host four invited talks by Thiago Castro Ferreira, Angela Fan, Behnam Hedayatnia and Lu Wang, as well as three non-archival presentations, and a panel on remaining challenges.

We would like to thank everyone who contributed to the success of this workshop, especially the authors, the program committee members, the organizers of the INLG 2019 conference and the INLG 2019 workshop chairs, and our sponsors, Saarland University and The Ohio State University.

—The Organizers

Organizers:

Anusha Balakrishnan, Microsoft Semantic Machines
Vera Demberg, Saarland University
Chandra Khatri, Uber AI
Abhinav Rastogi, Google AI
Donia Scott, University of Sussex / Scott Rush Associates
Marilyn Walker, University of California, Santa Cruz
Michael White, The Ohio State University / Facebook AI

Program Committee:

Paul Crook, Facebook AI
Alessandra Cervone, University of Trento
Claire Gardent, French National Centre for Scientific Research (CNRS)
Behnam Hedayatnia, Amazon Alexa AI
Dave Howcroft, Heriot-Watt University
Emiel Krahmer, Tilburg University
Shereen Oraby, University of California, Santa Cruz
Cecile Paris, CSIRO
Owen Rambow, Elemental Cognition
Alexander Rush, Harvard University
Frank Schilder, Thomson Reuters
Rajen Subba, Facebook AI
David Winer, University of Utah
Sam Wiseman, Toyota Technological Institute at Chicago
Amir Zeldes, Georgetown University
Yi-Chia Wang, Uber AI

Invited Speakers:

Thiago Castro Ferreira, Tilburg center for Cognition and Communication (TiCC), Tilburg University / Department of Linguistics, Federal University of Minas Gerais
Angela Fan, Facebook AI
Behnam Hedayatnia, Amazon Alexa AI
Lu Wang, CCIS, Northeastern University

Table of Contents

Invited Talks

<i>Data-to-text Natural Language Generation: Traditional, Novel and Future Methods</i> Thiago Castro Ferreira	I
<i>Convince Me If You Can: Argument Generation with Content Planning and Style Specification</i> Lu Wang	II
<i>Hierarchical Structure in Story Generation</i> Angela Fan	III
<i>Topical Chat: On the Structure of Knowledge Grounded Conversations</i> Behnam Hedayatnia	IV

DSNNLG Papers

<i>Maximizing Stylistic Control and Semantic Accuracy in NLG: Personality Variation and Discourse Contrast</i> Vrindavan Harrison, Lena Reed, Shereen Oraby and Marilyn Walker	1
<i>Incorporating Textual Evidence in Visual Storytelling</i> Tianyi Li and Sujian Li	13
<i>Fine-Grained Control of Sentence Segmentation and Entity Positioning in Neural NLG</i> Kritika Mehta, Raheel Qader, Cyril Labbe and François Portet	18
<i>Zero-shot Chinese Discourse Dependency Parsing via Cross-lingual Mapping</i> Yi Cheng and Sujian Li	24

Conference Program

10:00–11:30 Session 1

- 10:00–10:15 Welcome and Opening Remarks
- 10:15–11:00 *Data-to-text Natural Language Generation: Traditional, Novel and Future Methods* (Invited Talk)
Thiago Castro Ferreira
- 11:00–11:30 *Maximizing Stylistic Control and Semantic Accuracy in NLG: Personality Variation and Discourse Contrast*
Vrindavan Harrison, Lena Reed, Shereen Oraby and Marilyn Walker
- 11:30–12:00 Coffee Break

12:00–12:45 Session 2

- 12:00–12:45 *Convince Me If You Can: Argument Generation with Content Planning and Style Specification* (Invited Talk)
Lu Wang
- 12:45–2:15 Lunch

2:15–3:20 Session 3

- 2:15–3:00 *Hierarchical Structure in Story Generation* (Invited Talk)
Angela Fan
- 3:00–3:20 *Incorporating Textual Evidence in Visual Storytelling*
Tianyi Li and Sujian Li

3:20–4:15 Coffee Break and Poster Session

- 3:20–4:15 DSNLNG and Non-archival Posters
(continued next page)

4:15–5:45 Session 5

- 4:15–5:00 *Topical Chat: On the Structure of Knowledge Grounded Conversations* (Invited Talk)
Behnam Hedayatnia
- 5:00–5:45 Panel

Conference Program (cont.)

3:20–4:15 Coffee Break and Poster Session

DSNNLG Posters

Fine-Grained Control of Sentence Segmentation and Entity Positioning in Neural NLG

Kritika Mehta, Raheel Qader, Cyril Labbe and François Portet

Zero-shot Chinese Discourse Dependency Parsing via Cross-lingual Mapping

Yi Cheng and Sujian Li

Non-archival Posters

Towards a Scalable & Controllable Computational Solution for Document Planning

Craig Thomson, Ehud Reiter and Somayajulu Sripada

Modeling Conversation Context by Adapting Cognitive Architectures

Sashank Santhanam and Samira Shaikh

Constrained Decoding and Query Attention for Neural NLG in Task-Oriented Dialogue

Anusha Balakrishnan, Jinfeng Rao, Kartikeya Upasani, Michael White and Rajen Subba

Invited Talk

Data-to-text Natural Language Generation: Traditional, Novel and Future Methods

Thiago Castro Ferreira

Tilburg center for Cognition and Communication (TiCC), Tilburg University
Department of Linguistics, Federal University of Minas Gerais

Data-to-text Natural Language Generation (NLG) is a consolidated field of research which normally combines Computational Linguistics, Software Engineering and Artificial Intelligence methods to generate natural language from non-linguistic representations. Traditionally, most data-to-text applications have been designed using a modular pipeline architecture, in which the non-linguistic input data is converted into natural language through several intermediate transformations. In contrast, influenced by the phenomenon of deep learning, recent neural models for data-to-text generation have been proposed as end-to-end approaches, where the non-linguistic input is rendered in natural language with much less explicit intermediate representations in-between. Theoretically, we know that pipeline approaches are more transparent and their modules can be reused across applications, whereas neural end-to-end approaches may demand less manual labor and has registered state-of-the-art results in other text generation tasks, such as Machine Translation. Although we know these pros and cons in theory, the question about which kind of model empirically generates the most fluent and semantic texts from non-linguistic representations still remains unanswered. This lack of an empirical comparison is partially caused by the fact that traditional benchmarks for the task only consist of raw non-linguistic representations in parallel with their textual realizations, benefiting the evaluation of end-to-end approaches but not of pipeline ones, since explicit intermediate representations are missing for the study of particular modules of the latter architecture. In this presentation, I will introduce an annotation framework to enrich popular benchmarks with explicit intermediate representations, which will help the development and evaluation of particular pipeline modules as Discourse Ordering, Aggregation, Lexicalization, Referring Expressions Generation and Textual Realization. Based on a version of a popular data-to-text benchmark enriched with our framework, I will also present the results of a comparison between pipeline and end-to-end approaches. Finally, based on the findings of this project, I will project future challenges in the research of data-to-text NLG in terms of data, applications and evaluation.

Invited Talk

Convince Me If You Can: Argument Generation with Content Planning and Style Specification

Lu Wang
CCIS, Northeastern University

Understanding, evaluating, and generating arguments are crucial elements of the decision-making and reasoning process. A multitude of arguments and counter-arguments are constructed on a daily basis to persuade and inform us on a wide range of issues. However, constructing persuasive arguments is a challenging task for both human and computers, as it requires credible evidence, rigorous logical reasoning, and sometimes emotional appeals.

In this talk, I will introduce our neural network-based argument generation model. It consists of a powerful retrieval system and a novel two-step generation model, where a text planning decoder first decides on the main talking points and a proper language style for each sentence, then a content realization component constructs an informative and fluent paragraph-level argument. We believe that the proposed argument generation framework will enable many compelling applications, including providing unbiased perspectives on complex issues, debate coaching, and essay writing tutoring. Our framework is also generic and has been applied to other text generation problems, such as Wikipedia article paragraph generation and scientific paper abstract writing.

Invited Talk

Hierarchical Structure in Story Generation

Angela Fan
Facebook AI

We explore the task of story generation: creative systems that can build coherent and fluent passages of text about a topic. Using a dataset of 300k human-written stories paired with writing prompts, we investigate hierarchical story generation. Our models first generate a premise and then transform it into a passage of text. We develop models that improve the relevance of the story to the premise using a novel form of model fusion and present improvements to self-attention that better capture long-range context. Then, we build upon this work by proposing a coarse to fine mechanism for story generation, decomposing the task into several steps. We first explicitly generate logical verb sequences to model action in stories, then form these into sentences, and finally fill-in character names. We show that such decompositions improve the consistency and diversity of generated stories.

Invited Talk

Topical Chat: On the Structure of Knowledge Grounded Conversations

Behnam Hedayatnia

Amazon Alexa AI

Conversational agents like Amazon Alexa, Google Assistant and Apple Siri have been exploding in popularity over the past few years. However, much work remains in the area of social conversation over a broad range of domains and topics. To advance the state of the art in open domain dialog, Amazon launched the Alexa Prize, a 2.5-million-dollar university competition where selected university teams were challenged to build conversational agents, known as “socialbots”, to converse coherently and engagingly with humans on popular topics such as Sports, Politics, Entertainment, Fashion and Technology for 20 minutes. The Alexa Prize offers a unique opportunity to perform research and interact with real user conversational data at scale. Over the past two years, we have learned that there are certain areas that these bots could improve on such as topical depth, breadth and smooth topical transitions in order to have deep and engaging conversations. Given this information, we formed a conversational dataset where we can study how to create engaging conversations. We introduce Topical-Chat: a knowledge-grounded human-human conversation dataset, where the underlying knowledge spans 8 broad topics. Our dataset enables models to leverage world knowledge while conversing with humans leading to more coherent and interesting conversations. We will present some modeling work using generative encoder-decoder conversational models trained on Topical-Chat and perform automated and human evaluation for benchmarking. Additionally we will present an analysis of Topical-Chat based on the knowledge content selected and presented to humans as background knowledge.