

NLP for Conversations

Gabriel Murray

UFV

Canada

`gabriel.murray@ufv.ca`

Shafiq Joty

NTU

Singapore

`srjoty@ntu.edu.sg`

Giuseppe Carenini

UBC

Canada

`carenini@cs.ubc.ca`

Abstract

1 Tutorial Overview

The primary goal of this tutorial is for attendees to learn about recent work applying NLP to spoken and written conversations, with a focus on computational models for three related topics: conversational structure, summarization and sentiment detection, and group dynamics. We provide examples of specific NLP tasks within those three areas, how they relate to one another, their applications, and how we evaluate task performance.

We will begin by discussing motivations and applications of applying NLP methods to conversations, including downstream applications that could benefit. Attendees will hear about the challenges of working with noisy data, and examples of datasets of spoken and/or written conversations.

The first part of the tutorial covers **conversational structures**, the basic building blocks for working with conversational data. Participants will learn about computational methods for uncovering thread and topic structures of a conversation, detecting dialogue acts and adjacency pairs, identifying participant roles (where relevant), and how to treat disfluencies. We will cover methods for both synchronous (e.g., meeting, phone) and asynchronous (e.g., forum, email) conversations.

In the second part of the tutorial, we will focus on **sentiment analysis and summarization**. Attendees will learn about the related, overlapping tasks of detecting sentiment, subjectivity, and opinions. We will cover unsupervised and supervised approaches, as well as multimodal sentiment detection. Participants will learn about intrinsic vs. extrinsic evaluation of sentiment analysis methods for conversations.

For summarization, we will cover core topics, such as the notions of extractive vs. abstractive summarization, and summarization vs. compression. In particular, participants will learn about the limits of extractive summarization on noisy and opinion-filled conversation data. We will particularly emphasize the question of how to evaluate automatically generated summaries, including some of the controversial history surrounding automatic summarization metrics that are widely used.

In the final part of the tutorial, participants will learn about the growing field of research that uses NLP and machine learning methods to model and predict **group dynamics**, including prediction of group performance and participant affect. Attendees will learn about the close relationship between these three areas of summarization, sentiment, and group dynamics, and why researchers in each one of those areas often end up being concerned with the other two topics as well. Finally, we will discuss promising current and future directions of applying NLP to conversations.

2 Tutorial Outline

Introduction

- Types of conversational data (e.g., meetings, discussion forums, email)
- Challenges of conversational data
- Working with both spoken and written data

- Examples of conversational datasets (e.g., AMI, ELEA, BC3, Ubuntu Dialogs)
- Applications:
 - meeting support
 - social media analytics
 - business intelligence
 - people analytics / HR analytics
 - computational social science

Part 1: Building Blocks (Conversational Structures)

- Synchronous & asynchronous conversations
- Thread detection & conversation disentanglement
- Dialogue act & adjacency pairs
- Role detection
- Treating disfluencies

Part 2: Sentiment and Summarization

Sentiment

- Concepts and terminology: sentiment, subjectivity, polarity, and opinions
- Annotation schemes
- Annotated datasets (e.g., MPQA, AMI, BC3, Congressional floor debates, CMU-MOSI)
- Sentiment detection approaches:
 - lexicon-based
 - supervised approaches
 - unsupervised approaches
 - fine-grained vs. coarse-grained
 - multimodal sentiment detection
- Evaluation Metrics and Approaches (e.g., intrinsic vs. extrinsic)
- Applications
- Negative sentiment in face-to-face vs. online conversations

Summarization

- Concepts and terminology:
 - Abstraction vs. Extraction
 - Indicative vs. Informative
 - Summarization vs. Compression
- Summarization approaches:
 - Unsupervised extraction
 - Supervised extraction
 - NLG-based abstraction

- Hybrid methods
- Neural methods
- Evaluation Metrics and Approaches
 - ROUGE
 - Pyramids
 - Weighted precision-recall-fscore
 - Extrinsic evaluation
 - Human judgment, crowd-sourcing
- Annotated datasets (e.g., ICSI, AMI, BC3)
- The limits of extraction for conversational data
- The challenges of summarizing subjective data

Part 3: Group Dynamics

- Concepts and terminology: groups, dyads, task-centered groups
- Detecting emergent leadership from group discussion
- Measuring participant satisfaction
- Measuring group progress and performance
- Group cohesion
- Markov rewards analysis for group conversation data
- Social network analysis + NLP for small groups
- NLP and group conflict
- Datasets for analyzing speech/language and group interaction/performance (e.g. AMI, ELEA, Leadership corpus, Team Entrainment corpus)

3 Instructors

Dr. Gabriel Murray

Associate Professor
 Computer Information Systems
 University of the Fraser Valley
 Abbotsford, Canada
gabriel.murray@ufv.ca

Dr. Shafiq Joty

Assistant Professor
 Computer Science & Eng.
 Nanyang Technological University
 Singapore
srjoty@ntu.edu.sg

Dr. Giuseppe Carenini

Associate Professor
 Computer Science
 University of British Columbia
 Vancouver, Canada
carenini@cs.ubc.ca

Dr. Gabriel Murray is an Associate Professor in Computer Information Systems at the University of the Fraser Valley (UFV), and an Affiliate Professor in CS at University of British Columbia (UBC). His background is in computational linguistics and multimodal speech and language processing. He holds a PhD in Informatics from the University of Edinburgh, completed under the supervision of Drs. Steve Renals and Johanna Moore. His research has focused on various aspects of multimodal conversational data, including automatic summarization and sentiment detection for group discussions. Recent research also focuses on predicting group performance and participant affect in conversational data. In 2011, Gabriel co-authored the book “Methods for Mining and Summarizing Text Conversations.” He is the recipient of an NSERC Discovery grant on developing and applying machine learning methods for small group interaction.

Dr. Shafiq Joty is an Assistant Professor at the School of Computer Science and Engineering, NTU. Previously, he was a research scientist at the Qatar Computing Research Institute (QCRI). He holds a PhD in Computer Science from the University of British Columbia. His work has primarily focused on developing discourse analysis tools (e.g., discourse parser, coherence model, topic model, dialogue act recognizer), and exploiting these tools effectively in downstream NLP applications like machine translation, summarization, and sentiment analysis. Apart from discourse and its applications, he has also developed novel machine learning models for question answering, machine translation, and opinion analysis. Shafiq is a recipient of NSERC CGS-D scholarship and Microsoft Research Excellent Intern award.

Dr. Giuseppe Carenini is an Associate Professor in Computer Science at UBC. Giuseppe has broad interdisciplinary interests. His work on natural language processing and information visualization to support decision making has been published in over 100 peer-reviewed papers (including best paper at UMAP-14 and ACM-TiiS-14). Dr. Carenini was the area chair for “Sentiment Analysis, Opinion Mining, and Text Classification” of ACL 2009, the area chair for “Summarization and Generation” of NAACL 2012, the Program Co-Chair for IUI 2015, and the Program Co-Chair for SigDial 2016. He has also co-edited an ACM-TIST Special Issue on “Intelligent Visual Interfaces for Text Analysis.” In 2011, he published a co-authored book on “Methods for Mining and Summarizing Text Conversations” (?). In his work, Dr. Carenini has also extensively collaborated with industrial partners, including Microsoft and IBM. Giuseppe was awarded a Google Research Award, an IBM CASCON Best Exhibit Award, and a Yahoo Faculty Research Award in 2007, 2010 and 2016 respectively.