Detecting Discourse Roles in Online Dialogues: a Quantitative Approach

George Broadwell (University at Albany, US)
Umit Boz (University at Albany, US)
Tomek Strzalkowski (SUNY Albany, US)
Jennifer Stromer-Galley (University at Albany, US)
Samira Shaikh (University at Albany, US)
Sarah Taylor (Lockheed Martin, US)
Nick Webb (University at Albany, US)

Abstract:

This research is part of a larger project that involves developing computational tools to model and recognize communicative behavior in online environments. Specifically, the paper reports on a series of metrics which have been designed to reveal varying degrees of influence and involvement in online interactions.

Extended Abstract:

Measuring the quality of participant interaction in online discussions has been extensively studied in the field of online learning (Maor & Volet, 2007; Salmon, 2003; Garrison, Anderson and Archer, 2000). In this research, we report on a series of metrics which have been designed to reveal varying degrees of influence and involvement in online multi-party discourse. Our metrics incorporate a few well-known discourse measures, such as turn length and turn frequency, but also contain a number of novel measures of discourse behavior related to topic introduction, topic mention, and topic chains. We will demonstrate that these measures successfully model important dimensions of variability in participant behavior in such environments.

The overall objective of our research is to generate computational tools to model and recognize two kinds of communicative behavior (involvement and topic control) and associated social phenomena in online environments. Similar to Social Network Analysis (SNA), the automated tools will be used to identify the central and peripheral participants in an online multiparty dialogue by modeling the social phenomena of Power (Centrality) and Leadership based on the measures of participants' conversational behavior in discourse. In order to detect and attribute these social constructs to discourse participants, we developed a group of metrics which differentiate the participants according to their degree of involvement and influence in the discourse. Involvement in this research is defined as a degree of engagement or participation in online discussions. A degree of involvement may be estimated by how much a speaker contributes to the discourse in terms of substantive content. Highly involved speakers speak often and closely follow the conversation by making frequent utterances on topics that are relevant to the conversation. These speakers are particularly engaged in discussion of the most important and persistent local topics. Involvement in discourse is detected using the following indicators: a measure of information content in each participant's utterances (Noun Phrase Index), the frequency of turns per participant (Turn Index), the rate of participation in discussion of persistent local topics (Topic Chain Index), the rate of participation in discussion of local topics introduced by other speakers (All Subsequent Mentions and Allotopicality). Influence (Topic Control) refers to attempts by any discourse participants to impose the topic of conversation. Highly influential

participants introduce new topics into discourse, take up the topics introduced by others, and take sides on the topics being discussed. Similarly, a person who takes longer turns, whose topics are readily discussed by others, and to whom others tend to 'yield the floor', is considered to have a high degree of topic control in a conversation. The indicators of Topic Control in discourse are the rate of local topic introduction per speaker (Local Topic Introduction), rate of subsequent mentions of each topic per speaker (Subsequent Mentions), rate of subsequent mentions of local topics excluding self mentions (Cite-Score), and an average turn length per participant (Turn Length). In an online discourse that involved choosing between fictional candidates for a job, we enlisted 7 participants to engage in a 90 minute chat. These participants showed widely varied behavior in terms of their level of involvement in the discourse and in the extent to which they were able to influence the direction of the discourse. Looking at involvement metrics, some participants had as many as 31.3% (205/655) of all turns, while others had as few as 3.8% (25/655) of all turns. We capture this with the Turn Index (TI) metric. Other metrics related to involvement were highly correlated with the TI metric. For influence measures, we used a novel metric, Local Topic Introductions (LTI), which tracked which participants introduced the most local topics into the discourse, where we defined a local topic as any noun phrase introduced and subsequently referred to by a pronoun, synonym, or repetition. Again, participants varied widely in the number of local topics which they successfully introduced into the discourse, with numbers ranging from a high of 25.6% (23/90) of all topic introductions to a low of 1%(1/90) of all topic introductions. Our other metrics for influence reinforce and are correlated with LTI. Final influence assessments per participant are obtained by combining the predictions of all indicators (high, medium, low based on quintiles). Comparing results for influence and involvement in online discourse allows us to gain a clearer understanding of the differing roles that participants play in multiparty discourse, as the following abbreviated table of results indicates: Participant, Aggregate of Involvement metrics/Aggregate of Influence metrics; Comment JR, %25 (high)/%20 (medium); highly involved, moderately influential LE, %17 (medium)/%32 (high); moderately involved, highly influential KN, %16 (medium)/%21 (medium); moderately involved, moderately influential KI, %21 (high)/%18 (medium); highly involved, moderately influential CS, %5 (low)/%12 (low);less involved,less influential KA, %15 (medium)/%5 (low); moderately involved,less influential JY, %3 (low)/%2 (low); less involved,less influential

Because our project is developing a computational implementation of these metrics, the implication for online education is that it should be possible to use computational terms to show a profile of the differential linguistic behavior of a group of people engaged in online discourse. While our analysis, thus far, has been focused on synchronous online dialogues, the automated tools can also be used to investigate the discourse roles participants play in asynchronous online discussions. As an additional assessment tool to draw on, this would allow an educator to quickly see which participants are relatively uninvolved and/or uninfluential in various kinds of online discussions. When fully functional, the automated tools developed through this project will have far-reaching implications for research in online learning and social networking environments.

References:

Garrison, D. R., Anderson, T., & Archer, W. (2000). Critical inquiry in a text-based environment: Computer conferencing in higher education. The Internet and Higher Education, 2(2-3), 87-105.

Maor, D., & Volet, S. (2007). Interactivity in professional online learning: A review of research based studies. Australasian Journal of Educational Technology, 23(2), 269-290.

Salmon, G. (2003). E-moderating: The key to teaching and learning online. London: RoutledgeFarmer.