Modelling Student Participation Using Discussion Forum Data

Elaine Farrow
School of Informatics, University of Edinburgh, Edinburgh, UK
Elaine.Farrow@ed.ac.uk

ABSTRACT: Across many different educational settings, course discussion forums allow students to learn from one another and connect socially with their peers and instructors. Content analysis of the messages that are exchanged has been used to model engagement using two well-established theoretical frameworks, Community of Inquiry and ICAP. However, manual content analysis is slow and expensive, and prior work on automation is limited. In addition, these two theoretical frameworks developed out of different disciplines, and little work has been done to bring them together. To address these issues, I will evaluate the use of advanced methods from natural language processing to automate the content analysis, considering both frameworks individually and together, and comparing the results with prior work in terms of accuracy and explanatory power. I will also contribute to the conceptual understanding of what characterises a high quality discussion forum contribution by identifying connections between the frameworks themselves and places where they offer complementary perspectives.

Keywords: learning analytics, student engagement, Community of Inquiry, ICAP, natural language processing

1   INTRODUCTION

Course discussion forums are increasingly used to support large face-to-face classes, in addition to their on-going key role in online and distance learning courses. However, the volume of messages exchanged is often so great that instructors can struggle to read them all in a timely manner, or to identify common themes and threads of argument between them. These messages provide a rich source of material for researchers interested in studying how effective learning takes place through discussion (Garrison, Anderson, & Archer, 1999), and there is growing interest in using this data to create models of student engagement. Content analysis techniques can be used to classify the depth and quality of messages using labels from an educational framework, in order to identify conversation threads that are developing appropriately and those that have stalled or are off-task. Two popular frameworks for modelling student engagement are the Community of Inquiry (CoI) framework (Garrison, Anderson, & Archer, 1999), and the ICAP framework (Chi & Wylie, 2014). CoI is one of the best-studied theoretical frameworks in online education (Gašević, Adesope, Joksimović, & Kovanović, 2015), and ICAP has been used as a foundation for many studies on computer-supported collaborative learning (Wang, Yang, Wen, Koedinger, & Rosé, 2015). By automating the content analysis, the results can be used while a course is still running. For example, instructors could be notified about conversation threads where they might want to intervene (although the specifics of that intervention are out of scope for this research project). Automation also allows research to be done on large data sets where manual annotation is impractical. Computational models that can assign labels to new data can also provide further insights by revealing patterns within the data. For example, a random forest model can report which of the model features was most informative.

In my doctoral work, I will make use of both the CoI and ICAP frameworks and evaluate the ability of advanced methods from Natural Language Processing (NLP) to automate the content analysis based on the labelling schemes provided by the two frameworks. My goal is to improve the way we identify and model the depth and quality of student participation using discussion forum data. I aim to develop methods that handle input text more flexibly, while producing outputs that are at least as accurate and informative as previous work. My work will also contribute to a better conceptual understanding of engagement through analysis of the relationship between the frameworks.

2   BACKGROUND

2.1   Theoretical frameworks for modelling student engagement

2.1.1   Community of Inquiry (CoI)

The Community of Inquiry (CoI) framework for online education is a powerful tool for analysing and developing effective learning experiences (Garrison, Anderson, & Archer, 1999). The framework identifies three main elements ('presences') that are important for a successful educational experience: i) a social environment conducive to learning (*social presence*); ii) a well-designed course with on-going facilitation (*teaching presence*); and iii) the student's own cognitive engagement with the subject matter (*cognitive presence*). CoI has been widely used to analyse student learning in online courses (Gašević, Adesope, Joksimović, & Kovanović, 2015), and predictive models have been developed for identifying its elements automatically using the text of discussion forum messages (e.g., Waters, Kovanović, Kitto, & Gašević, 2015).

Two recent studies (Kovanović, et al., 2016; Neto, et al., 2018) that developed models for predicting the phases of cognitive presence both reported high accuracy scores for the prediction task, using linguistically-motivated features as input to the model – things like text coherence, complexity, and readability scores. These were derived from the messages using the text analysis tools LIWC (Linguistic Inquiry and Word Count) (Tausczik & Pennebaker, 2010) and Coh-Metrix (McNamara, Graesser, McCarthy, & Cai, 2014). The features were chosen because they have potential explanatory power, and the studies explored which of them were most predictive. However, the value of the feature analysis is called into question by doubts surrounding the validity of the models themselves. A replication study (Farrow, Moore, & Gašević, 2019) showed that data contamination between the training and testing data in these studies could have led to over-optimistic accuracy scores. Furthermore, only 9 of the top 20 most predictive features from one study (Kovanović, et al., 2016) were still in the top 20 after avoiding the potential contamination, suggesting that over-fitting may have led the prior model to see some features as more predictive than was really the case, and to disregard others that actually have more discriminative power (Farrow, 2018). Therefore, further investigation is needed into the features that characterise high quality discussion contributions.

2.1.2   ICAP

The ICAP framework (Chi & Wylie, 2014) defines cognitive engagement based on overt behaviours alone. The framework looks at individual learning activities and how they relate to students' cognitive engagement with the learning materials. Four 'modes' of engagement are identified, and the framework predicts that higher modes will be correlated with greater learning gains. The four modes, in descending order, are **I**nteractive, **C**onstructive, **A**ctive, and **P**assive. Each of these modes subsumes the modes below it and represents a qualitatively different *kind* of growth in knowledge,

not simply a bigger change. Passive engagement corresponds to the least taxing on-task activities; for example, listening to a lecture. Active engagement covers activities that demand the student's attention, such as highlighting lecture notes. To qualify as constructive engagement, novel output must be generated – for example, summary notes. Interactive engagement requires interaction with a partner, and normally both partners must be engaged constructively. However, this requirement is relaxed in the case of activities involving larger groups, since subsets of participants can engage with the same task in different ways. Off-task behaviours do not constitute any type of engagement. ICAP has recently been used to classify student comments on MOOC videos (Taskin, Hecking, Hoppe, Dimitrova, & Mitrovic, 2019). Modified versions of ICAP have been used to analyse discussion forum messages in MOOCs (Wang, Wen, & Rosé, 2016) and student comments on an annotated electronic course text (Yogev, Gal, Karger, Facciotti, & Igo, 2018). Future work can build on this foundation.

2.1.3   Comparing the frameworks

While both frameworks address engagement, they do so from different perspectives. They were developed independently and with different goals in mind. CoI was developed specifically in order to understand the benefit of online education and to explain how students are able to learn and develop ideas through discussion. ICAP has a broader scope and has been demonstrated to be effective in predicting the educational value of several different interventions, in a classroom setting as well as online. Little prior work has been done to compare the frameworks, either conceptually or through experimentation. If the labels they assign to messages are found to be closely correlated, then results derived using each of them in previous studies can be expected to be applicable to work using the other. If, instead, they are completely distinct, then using them together in future studies will give a richer picture of engagement. A triangulation study involving both conceptual and empirical comparisons of the frameworks would thus offer a useful contribution to the theoretical understanding of online learning, critical discourse, and learning through discussion.

2.2   Neural network models and advanced NLP methods

In recent years, the field of natural language processing has increasingly embraced the use of neural network methods to classify text automatically. State-of-the-art neural networks can be used to produce accurate outputs for many application domains using only text as input, without the need for extensive feature engineering (Goodfellow, Bengio, & Courville, 2016). Many such applications make use of pre-trained language models known as **word embeddings** (Mikolov, Sutskever, Chen, Corrado, & Dean, 2013; Pennington, Socher, & Manning, 2014), which transform words into points in a high-dimensional vector space. In the high-dimensional space, words with similar meanings are found near one another, while dissimilar words are far apart. This means synonyms are treated in similar ways and common spelling errors can be handled automatically. This approach could be very beneficial for automated content analysis of forum messages, which often contain misspellings.

Neural network models can be hard to interpret, but the use of an **attention layer** (Wu, et al., 2016) is often described as allowing researchers to 'see inside' what is otherwise a black-box technique. After transforming each of the input words into a vector and processing those vectors through the early layers of the network, the attention layer combines the results in a weighted sum before passing them on to later layers. The learned weights in the attention layer can thus be seen as indicating the strength of influence of each of the input words on the final output classification; that is, the extent to which each of the words in a message determines the quality of the contribution.

Work on **multi-task and transfer learning** (Collobert & Weston, 2008) has shown that training a single neural network to learn to generate multiple target outputs at the same time can help to avoid over-fitting to the training data and produce better models overall. This suggests that learning the labels for both CoI and ICAP together could work better than using either framework alone.

3   AIMS OF THE RESEARCH

The overall problem that my research will address is how we can identify and model the depth and quality of student participation using the messages that students post to course discussion forums. My research has two main goals: 1) to discover where the CoI and ICAP frameworks take a similar approach and where they provide complementary insights; and 2) to evaluate the use of advanced NLP methods to automate the labelling process on new data. Specifically, I will investigate the performance of models that use techniques including word embeddings, attention layers, and multi-task and transfer learning. This work aims to answer four specific research questions.

RQ1: What is the association between the phases of cognitive presence in the CoI framework and the modes of engagement in the ICAP framework?

RQ2: If pre-trained language models such as word embeddings are used to automate message labelling, is model performance comparable with prior studies that used linguistically motivated features to train the model?

RQ3: Can an attention layer in a neural network reveal what aspects of a discussion forum message are important for identifying depth and quality of participation?

RQ4: Does model performance improve when labels from CoI and ICAP are learned together, compared to the performance of models using each framework separately?

4   METHODOLOGY

My research combines methodological work with quantitative modelling and qualitative content analysis. My current study (target date for completion: early 2020) will compare the CoI and ICAP frameworks by looking at co-occurrences of ICAP modes with phases of cognitive presence in a manually labelled data set. This is anonymised data that was collected in a previous study and ethical approval has already been obtained. Specifically, I will approach this task quantitatively by looking at confusion matrices between labels from the two frameworks and visualising them using Epistemic Network Analysis (ENA) (Shaffer, Collier, & Ruis, 2016), as well as comparing the frameworks theoretically and conceptually (**RQ1**). My expectation is that the two frameworks are sufficiently distinct that they will provide complementary insights into the learning processes demonstrated in discussion forum messages.

Later studies will look at automating the labelling of discussion forum messages using advanced NLP methods. I will develop neural network models that incorporate word embeddings and an attention layer (target date for completion: April 2020) and compare the performance of these models with simpler predictive models such as random forests – both quantitatively, in terms of model performance (**RQ2**), and also qualitatively, in terms of potential explanatory power (**RQ3**). By mapping the words into a high-dimensional vector space using word embeddings, the effects of

particular word choices are expected to diminish. Therefore, I expect that this approach could prove to be just as powerful as using linguistically motivated model features, while adding flexibility. An attention layer could indicate which words and phrases best characterise the depth and quality of participation according to each of the theoretical frameworks. These results can be validated qualitatively by comparison with prior work on factors contributing to student engagement. One potential future application of this aspect of the research could perhaps be the automatic generation of hints for students about how to improve their own discussion contributions. Finally, I will use multi-task and transfer learning to train models using the labels from both frameworks at once (target date for completion: July 2020) and compare their performance with models trained on each set of labels individually, addressing **RQ4**. If performance improves, in line with prior work, this result would also support the use of both frameworks together in future studies.

5   CURRENT STATUS AND RESULTS ACHIEVED

My first methodological study looking at how data contamination can arise from common data pre-processing practices was presented at LAK'19. A summary of this work was also shared with a broad data science audience at a UK-wide workshop (Advances in Data Science 2019). I am now working on data preparation for my next experimental study. I have adapted the extended ICAP coding scheme used in prior work (Yogev, Gal, Karger, Facciotti, & Igo, 2018) to be more relevant to the context of the data set that I am using. The messages were already labelled with phases of cognitive presence, and manual annotation with the labels from my adapted ICAP scheme is in progress. A study based on preliminary analysis of this data was accepted as a short paper in the main LAK'20 research track.

REFERENCES

Chi, M. T., & Wylie, R. (2014). The ICAP Framework: Linking Cognitive Engagement to Active Learning Outcomes. *Educational Psychologist, 49*(4), 219-243. doi:https://doi.org/10.1080/00461520.2014.965823

Collobert, R., & Weston, J. (2008). A unified architecture for natural language processing. *Proceedings of the 25th International Conference on Machine Learning* (pp. 160-167). New York, NY, USA: ACM. doi:https://doi.org/10.1145/1390156.1390177

Farrow, E. (2018). *Automated content analysis for modelling student engagement in online discussion forums.* Master's Thesis, School of Informatics. University of Edinburgh.

Farrow, E., Moore, J., & Gašević, D. (2019). Analysing discussion forum data: a replication study avoiding data contamination. *Proceedings of the 9th International Learning Analytics & Knowledge Conference (LAK19)*, (pp. 1-10). doi:https://doi.org/10.1145/3303772.3303779

Garrison, D. R., Anderson, T., & Archer, W. (1999). Critical Inquiry in a Text-Based Environment: Computer Conferencing in Higher Education. *The Internet and Higher Education, 2*(2), 87-105. doi:https://doi.org/10.1016/S1096-7516(00)00016-6

Gašević, D., Adesope, O., Joksimović, S., & Kovanović, V. (2015). Externally-facilitated regulation scaffolding and role assignment to develop cognitive presence in asynchronous online discussions. *The Internet and Higher Education, 24*, 53-65. doi:https://doi.org/10.1016/j.iheduc.2014.09.006

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning.* MIT Press.

Kovanović, V., Joksimović, S., Waters, Z., Gašević, D., Kitto, K., Hatala, M., & Siemens, G. (2016). Towards Automated Content Analysis of Discussion Transcripts: A Cognitive Presence Case.

Proceedings of the Sixth International Conference on Learning Analytics & Knowledge (pp. 15-24). New York, NY, USA: ACM. doi:https://doi.org/10.1145/2883851.2883950

McNamara, D. S., Graesser, A. C., McCarthy, P. M., & Cai, Z. (2014). *Automated Evaluation of Text and Discourse with Coh-Metrix.* Cambridge: Cambridge University Press. doi:https://doi.org/10.1017/CBO9780511894664

Mikolov, T., Sutskever, I., Chen, K., Corrado, G., & Dean, J. (2013). Distributed Representations of Words and Phrases and their Compositionality. *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2* (pp. 3111-3119). Curran Associates, Inc.

Neto, V., Rolim, V., Ferreira, R., Kovanović, V., Gašević, D., Dueire Lins, R., & Lins, R. (2018). Automated Analysis of Cognitive Presence in Online Discussions Written in Portuguese. *Proceedings of the 13th European Conference on Technology Enhanced Learning* (pp. 245-261). Cham: Springer International Publishing. doi:https://doi.org/10.1007/978-3-319-98572-5_19

Pennington, J., Socher, R., & Manning, C. (2014). GloVe: Global Vectors for Word Representation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 1532-1543). Association for Computational Linguistics. doi:https://doi.org/10.3115/v1/D14-1162

Shaffer, D. W., Collier, W., & Ruis, A. R. (2016). A Tutorial on Epistemic Network Analysis: Analyzing the Structure of Connections in Cognitive, Social, and Interaction Data. *Journal of Learning Analytics, 3*(3), 9-45. doi:https://doi.org/10.18608/jla.2016.33.3

Taskin, Y., Hecking, T., Hoppe, H. U., Dimitrova, V., & Mitrovic, A. (2019). Characterizing Comment Types and Levels of Engagement in Video-Based Learning as a Basis for Adaptive Nudging. *Proceedings of the 14th European Conference on Technology Enhanced Learning* (pp. 362-376). Cham: Springer International Publishing.

Tausczik, Y. R., & Pennebaker, J. W. (2010). The Psychological Meaning of Words: LIWC and Computerized Text Analysis Methods. *Journal of Language and Social Psychology, 29*(1), 24-54. doi:https://doi.org/10.1177/0261927X09351676

Wang, X., Wen, M., & Rosé, C. P. (2016). Towards Triggering Higher-order Thinking Behaviors in MOOCs. *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge* (pp. 398-407). New York, NY, USA: ACM. doi:https://doi.org/10.1145/2883851.2883964

Wang, X., Yang, D., Wen, M., Koedinger, K., & Rosé, C. P. (2015). Investigating how student's cognitive behavior in MOOC discussion forums affect learning gains. *Proceedings of the 8th International Conference on Educational Data Mining* (pp. 226-233). International Educational Data Mining Society (IEDMS).

Waters, Z., Kovanović, V., Kitto, K., & Gašević, D. (2015). Structure Matters: Adoption of Structured Classification Approach in the Context of Cognitive Presence Classification. *Information Retrieval Technology, 9460*, 227-238. doi:https://doi.org/10.1007/978-3-319-28940-3_18

Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., . . . Dean, J. (2016). Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation. *CoRR*. Retrieved from http://arxiv.org/abs/1609.08144

Yogev, E., Gal, K., Karger, D., Facciotti, M. T., & Igo, M. (2018). Classifying and Visualizing Students' Cognitive Engagement in Course Readings. *Proceedings of the Fifth Annual ACM Conference on Learning at Scale* (pp. 1-10). New York, NY, USA: ACM. doi:10.1145/3231644.3231648