A network analytic approach to integrating multiple quality measures for asynchronous online discussions

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

Dragan Gašević
Dragan.Gasevic@monash.edu
Faculty of Information Technology, Monash University
Clayton 3800, Australia

ABSTRACT

Asynchronous online discussions within a community of learners can improve learning outcomes through social knowledge construction, but the depth and quality of student contributions often varies widely. Approaches to assessing critical discourse typically use content analysis to identify indicators that correspond to framework constructs, that in turn serve as measures of depth and quality. Often only a single construct is addressed for performing content analysis in the literature, although recent work has used both social presence and cognitive presence constructs from the Community of Inquiry (CoI) framework. Nevertheless, there is no effective, commonly used, analytic approach to combining insights from multiple perspectives about quality and depth of online discussions. This paper addresses the gap by proposing the combined use of cognitive engagement (the ICAP framework) and cognitive presence (CoI); and by proposing a network analytic approach that quantifies the associations between the two frameworks and measures the moderation effects of two instructional interventions on those associations. The present study found that these associations were moderated by one intervention but not the other; and that messages labelled with the most common phase of cognitive presence could be usefully assigned to smaller meaningful subgroups by also considering the mode of cognitive engagement.

CCS CONCEPTS

• **Computing methodologies** → **Modeling methodologies**; • **Applied computing** → *Collaborative learning*.

KEYWORDS

discussion forum, critical thinking, Community of Inquiry, cognitive presence, ICAP, Epistemic Network Analysis

ACM Reference Format:

Elaine Farrow, Johanna Moore, and Dragan Gašević. 2021. A network analytic approach to integrating multiple quality measures for asynchronous online discussions. In *LAK21: 11th International Learning Analytics and Knowledge Conference (LAK21), April 12–16, 2021, Irvine, CA, USA.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3448139.3448163

LAK21, April 12–16, 2021, Irvine, CA, USA

ACM ISBN 978-1-4503-8935-8/21/04…$15.00
https://doi.org/10.1145/3448139.3448163

1 INTRODUCTION

Text-based online discussion forums have been used in a wide variety of educational settings for many years [14, 35]. In the context of the recent sudden and substantial increase in online education brought about by a global pandemic, research into the learning that takes place as students interact in such a setting is both timely and important. The very act of composing and submitting messages to a forum of peers can directly lead to learning, particularly when it involves refining and negotiating meaning in cooperation with others [10]. Previous studies found that participation in discussion forums was positively correlated with learning gains in MOOCs, even in the cases where messages received no reply or only a superficial response [31, 32, 34]. This could be attributed to the known benefits of self-explanation [4], which does not require a conversational partner. However, both the 'private world' of reflection and the 'shared world' of discussion play a vital role in learning, according to the practical inquiry model [15, 17].

This study aims to develop a richer representation of the depth and quality of student participation in online discussion forums. In this, we follow Schrire in taking "a perspective emphasizing the depth and quality of learning rather than the quantity" [28, p. 54]. Specifically, this study brings together insights from two of the most widely-used and well-supported frameworks. We posit that the frameworks are related, yet different in important ways; and, therefore, that they provide complementary perspectives to the study of the depth and quality of participation. The Community of Inquiry framework (CoI) [16] was designed specifically to support online learning through computer-mediated discussion. Meanwhile, the ICAP framework [5] has been applied successfully in many different educational situations – classroom-based as well as online. Previous work has considered the link between certain dialogue attributes and measures of participation in these two frameworks [8], but did not look directly at how the constructs were related to one another. Analysing the correspondences between constructs from the two frameworks would allow us to develop a richer approach to the measurement of depth and quality in online discussions than using either one alone. In this work, our specific intent is to inform the development of a learning analytic approach to measuring the association between the frameworks and assessing whether that association is moderated by instructional scaffolds, given the importance of instructional guidance to support social knowledge construction in online discussions [13].

In the current study, we focused on the definitions of the *phases of cognitive presence* from CoI and the *modes of cognitive engagement* in ICAP. We contribute to the existing body of literature on assessing the depth and quality of student participation in online discussions in the following ways: 1) We theorise about the general associations between the constructs of two widely-used theoretical frameworks, grounding our arguments in empirical data. 2) We propose a novel network analytic approach to the analysis of associations between the two theoretical frameworks and demonstrate how it can be used in assessment of depth and quality of online discussions. 3) We examine the associations between the phases of cognitive presence and the modes of cognitive engagement and how they are moderated by instructional scaffolds aimed at promoting cognitive presence.

2 BACKGROUND

2.1 Analysis of social knowledge construction in online discussion

The analysis of social knowledge construction through collaborative asynchronous discussion can be investigated from multiple perspectives and at different levels of granularity: from coarse-grained patterns of interaction between participants, through message-based cognitive and social presence labels, to fine-grained discourse analysis looking at the individual dialogue acts and conversational moves within a message [26, 28]. The content of individual messages can be labelled, through content analysis [7], using various theoretical frameworks to identify expressions of critical thinking and conversational moves that support the establishment of a productive and supportive online community. Two of the most widely-used frameworks are the Community of Inquiry framework (CoI) [15, 16] and the ICAP framework [5], discussed in more detail below (Sections 2.2 and 2.3). Both frameworks define coding schemes that allow annotators to assign framework labels consistently to previously unseen content, for the purposes of content analysis. Recent work has identified cue phrases and dialogue features that are correlated with specific framework labels [8, 10], and relationships between the labels and topics extracted from the course content [27]. Although manual content analysis is slow and expensive, several recent studies have achieved promising results automating the labelling process using these frameworks [1, 9, 11, 20, 21, 23, 25].

Combining multiple perspectives on the analysis of online discussions can bring further benefits. One study [28] that used a multi-framework approach found correspondences between the patterns of interactions among participants, the phases of critical inquiry, and the conversational moves in the messages. In the subset of discussion threads that were characterised by direct interactions between students, more of the messages were labelled as belonging to higher phases of critical thinking, compared with threads where most messages were responses to the instructor's original discussion prompt. Three different frameworks were used to assess the presence and extent of critical thinking. These were Bloom's Taxonomy [3], the SOLO taxonomy [2], and CoI [15, 16], referred to in that work as the Practical Inquiry Model. Schrire acknowledged that triangulating findings across the three models is "conceptually problematic since each model is based on a different theory of cognitive activity" [28, p. 66]. Nevertheless, broad correspondences were confirmed, supporting social constructivist learning approaches. Whereas previous work [28, 31] looked in general terms at indicators of high and low levels of critical thinking, the present study looked into the detail of the individual constructs in each framework and examined how the associations between them were moderated by two instructional interventions.

2.2 The Community of Inquiry framework

The Community of Inquiry (CoI) framework identifies three 'presences' that are important for successful learning: *social presence*, through which the members of a community of learners connect with one another as 'real people' [16]; *teaching presence*, including the design of the course and assignments as well as ongoing facilitation; and *cognitive presence*, which relates to the students' intellectual engagement with the course content.

Cognitive presence is considered to be the element most basic to educational success – in fact, both social and teaching presence function primarily as a support for cognitive presence [16]. It aims to capture the process whereby discussion participants construct meaning through communication. Cognitive presence can be conceptualised as a measure of the depth and quality of student participation [8]. It has four phases:

- **Triggering Event:** a question, task, or problem that triggers the process of critical inquiry.
- **Exploration:** an exploration of the task or problem through the exchange of ideas, but lacking selectivity.
- **Integration:** an examination and integration of ideas through identifying connections and constructing meaning.
- **Resolution:** a resolution of the original problem, coupled with building consensus among participants.

Progression through the phases is expected to develop over time as the activities in each phase build on the previous one, although progress is seldom linear. The topic of the discussion is set by the initial *Triggering Event*. Messages in the *Exploration* phase tend to be wide-ranging, though always linked back to the original topic. The integration of ideas and the beginnings of a coherent line of reasoning signal that the discussion has entered the *Integration* phase. Many discussions will not reach the *Resolution* phase without intervention from an instructor, since students often feel more comfortable in the earlier phases [17]. Resolution can involve consensus building within the community and the generation of possible solutions to the original problem, sometimes in the form of thought experiments. Clarification questions and new lines of inquiry can appear at any point in the discussion, perhaps beginning a new cycle. Messages containing indicators from multiple phases are coded with the highest phase (*coding up*) [33].

2.3 The ICAP framework

Cognitive engagement is the central concept in the ICAP framework [5]. The framework defines four modes of cognitive engagement, based on observable student behaviours: **I**nteractive, **C**onstructive, **A**ctive, and **P**assive.

- **Interactive:** an interaction with a partner while both are engaged constructively.
- **Constructive:** the generation of novel content; for example, through reasoning or summarisation.

- **Active:** an activity that demands attention, such as referencing previously given information.
- **Passive:** an activity where attention can wander; for example, reading without responding.

Higher modes are theorised to correlate with greater learning gains. Each mode represents a qualitative shift in knowledge growth. The original framework was adapted and expanded in several recent studies [8, 31, 36]. The label definitions for *Constructive* and *Active* were subdivided to introduce *Constructive Reasoning* and *Constructive Extending*, and *Active Targeted* and *Active General*, respectively. Messages consisting primarily of agreement or thanks were also given their own label, *Affirmation* [8]. In the context of online discussions, the *Interactive* label was reserved for messages that contained a direct response to the content of an earlier contribution. Responses to external information sources, such as textbooks or video presentations were instead labelled as *Constructive*.

The extended ICAP schema was used to label messages in MOOC discussion forums [31, 32]. It was also used to classify student comments on an annotated electronic course text [36] and on MOOC videos [30], although in the latter case the activity was carried out individually with no scope for interaction between participants. Results from these studies indicated that interaction was rare and the majority of messages did not receive contentful responses. The students who contributed the most frequently tended to generate questions rather than building on what others had written. An intervention to encourage students to generate more content [30] simply increased the proportion of shallow and simple (*Active* mode) messages and did not improve learning gains.

3 RESEARCH QUESTIONS

Section 2.1 presented the benefits of bringing together a variety of perspectives in order to achieve deeper understanding of how social knowledge construction takes place through asynchronous discussion. Previous studies achieved new insights by utilising both coarse-grained and fine-grained analyses [28] and by looking at the relationship between different presences within CoI [26]. There has been less research bringing together insights from multiple theoretical frameworks. In particular, the existing research has little to say about how the two most widely-used frameworks for assessing the depth and quality of student participation in online discussion forums are related. The first research question in the present study was therefore:

RQ1: *What are the general associations between the individual CoI phases of cognitive presence and the ICAP modes of cognitive engagement?*

The primary role of teaching presence is to support the development of cognitive presence [15, 16], so it is important to evaluate the effect of different instructional scaffolds on the development of concepts from the two frameworks and the associations between them. It is easy to imagine how changes in knowledge construction prompted by an intervention could improve one metric, while leaving another unaffected or even reduced – for example, by increasing the quantity of simple questions asked [30]. In this study, we specifically focused on two types of intervention to examine whether and how they moderated the associations between CoI and ICAP. One intervention involved the assignment of specific roles within the discussion, which is a common scaffolding approach in computer-supported collaborative learning [13]. The other one was based on increasing the scaffolding support to encourage deeper and higher quality engagement through providing quality standards that quality contributions to online discussion should have. Prior work showed that both of these interventions were effective in improving measures based on the CoI phases of cognitive presence [18], and that they had a moderating effect on the associations between cognitive presence and social presence [26], and between social presence and learning outcomes [22]. However, no analysis had looked at the effect on the ICAP modes of cognitive engagement, or any changes that might be seen in the patterns of connection between those and the CoI phases of cognitive presence. Thus, our second research question was:

RQ2: *Are, and if so, how are the associations between the CoI phases of cognitive presence and the ICAP modes of cognitive engagement moderated by the instructional scaffolds aimed at promoting cognitive presence?*

4 METHOD

4.1 Description of the data

The course discussion forum messages used in this study were collected during six sessions of the same online Masters level course at a Canadian university between 2008 and 2011. Each message thread corresponded to a discussion that was led by a single student, who took on the role of *Research Expert* for the duration of the thread. The opening post in each thread introduced the new topic and included a link to a video presentation the student had prepared based on their own choice of research paper. The other students in the thread took the *Practising Researcher* role, giving feedback on the presentation and asking questions about the content. The course instructors rarely got involved with the discussions – only 6 posts in total across the whole data set of 1,747 messages were written by an instructor. Forum participation was graded and accounted for 10% of the overall course mark.

4.2 Framework labels from CoI and ICAP

The messages were labelled with both the CoI phases of cognitive presence and the ICAP modes of cognitive engagement. Exactly one label from each framework was used for each message. If a message matched the description for more than one label in either framework, the label that corresponded to greater depth and quality of participation was chosen. The additional labels *Other* and *Off-task* were used where there were no indications of cognitive presence or cognitive engagement, respectively. They are not generally considered to be part of the frameworks.

The labels for the CoI phases of cognitive presence were assigned by two expert coders who achieved high levels of agreement (98.1% agreement, Cohen's κ = 0.974). The label distribution is shown in Table 1. In a similar way, the cognitive engagement labels of ICAP were assigned by two postgraduate students working independently, achieving 'substantial' inter-annotator agreement (Cohen's κ = 0.623) [24]. An expanded set of ICAP labels was used, following earlier work [8, 31, 32, 36] (see Section 2.3). The distribution of these labels across the data set is presented in Table 2.

Table 1: Breakdown of messages by CoI phases of cognitive presence in the whole data set. The *Other* label was used where a message did not display any cognitive presence.

Cognitive presence phase	Short label	Example behaviour	Count	Percentage
Triggering Event	**TRIGGERING**	Asking a question or posing a problem	308	17.63%
Exploration	**EXPLORATION**	Exchanging ideas	684	39.15%
Integration	**INTEGRATION**	Integrating ideas and constructing meaning	508	29.08%
Resolution	**RESOLUTION**	Reaching consensus or suggesting a new hypothesis	107	6.12%
Other	**OTHER**	Commenting with no signs of cognitive presence	140	8.01%
All			1,747	100.00%

Table 2: Distribution of cognitive engagement modes in the whole data set. The *Off-task* label was used for messages displaying no cognitive engagement. The *Passive* label was not used in this data set.

Cognitive engagement mode	Short label	Example behaviour	Count	Percentage
Interactive	**I**	As for **C1**, but in direct response to an earlier message	579	33.14%
Constructive Reasoning	**C1**	Displaying explanation or reasoning about the current topic	313	17.92%
Constructive Extending	**C2**	Introducing new content to the discussion	409	23.41%
Active Targeted	**A1**	Referencing specific previous content	75	4.29%
Active General	**A2**	Showing other signs of being engaged with course content	287	16.43%
Passive	**P**	Reading messages without responding	0	0.00%
Affirmation	**F**	Affirming what was said in an earlier message	73	4.18%
Off-task	**O**	Commenting with no relation to the current topic or the course	11	0.63%
All			1,747	100.00%

4.3 Scaffolding interventions

Two distinct scaffolding interventions took place during this course. One was the within-subjects role assignment described in Section 4.1, where each student took a turn at being the *Research Expert* and presented their work to a group of *Practising Researcher*s. The second intervention was a change in the participation instructions between the second and third sessions of the course, with additional guidance given about what would constitute a high quality contribution to the discussion (Figure 1). The first two sessions of the course thus constitute the *Control* group, while the remaining four sessions, where the additional guidance was given, are the *Treatment* group. The number of participants and messages exchanged in each instructional condition is shown in Table 3.

Table 3: Messages sent and unique participants in each instructional condition. Students were expected to take on both the *Research Expert* and *Practising Researcher* roles. The instructor and one student who repeated the course participated in both *Control* and *Treatment* groups.

	Control	Treatment	Total	Participants
Practising Researcher	434	466	900	81
Research Expert	411	436	847	82
Total	845	902	1,747	
Participants	39	48		85

Your participation needs to be about the content being presenting with the following three levels (from the lowest to the highest quality):

(i) clarification question – asking about some uncertain parts of the paper being presented;
(ii) synthesis question – asking a question that connects the topics of the presentation at hand with another peer-reviewed paper and its results covered either in the study guide, presentation of another student, or a peer-reviewed research publication;
(iii) innovation question – asking or proposing a novel research topic by making use of the results presented in the paper at hand to draw ideas that are formulating a research problem/challenge. Preferably, the result of a discussion triggered by such a question might result even in the problem formulation of the research to be done in the final assignment of the course.

Figure 1: Additional assignment instructions issued to the *Treatment* group.

4.4 Data analysis methods

4.4.1 Cross-tabulation. As a first step towards answering *each of our research questions*, we used a cross-tabulation approach to examine the distributions of the framework labels across the data set, in a similar manner to prior work [28]. Plots such as bar charts and heat maps are useful for visualising the results of cross-tabulation, in order to obtain a preliminary understanding of the data set as a whole, relevant to **RQ1**. They also allow for comparing different subsets, relevant to **RQ2**.

4.4.2 Epistemic Network Analysis. The second stage of our analysis used Epistemic Network Analysis (ENA) [29]. ENA is a network analytic approach that is ideal for investigating the relationships between relatively small sets of concepts across a densely connected network of interactions. It has been used successfully in many previous studies of cognitive presence and social presence [12, 19, 26, 27]. Different sub groups of analysis units (for example, students or messages) can be compared both visually and statistically, providing both qualitative and quantitative insights.

ENA is based around measuring and visualising the *connections between concepts*, rather than treating the concepts in isolation. The abstract high-dimensional graph of connections between concepts is projected down onto a lower-dimensional *projection space* using *singular value decomposition* (SVD). Network edges are weighted by the frequency of connections between pairs of concepts. This method ensures that data points with similar patterns of connections will generally appear close together in the projection space, while those that differ more will appear further apart. Additionally, the positioning of the concepts as nodes within the projection space can indicate which concepts tend to be linked to other concepts in similar ways, and thus can make the space itself interpretable [29].

When the network is projected down to the two most informative dimensions (that is, into a flat 2D representation), it is possible to visualise the network of connections for a single data point – for example, one student. Each individual network can also be summarised by a single point in the projection space that represents the *centroid*, or weighted mean, of that network. This allows comparisons on a larger scale, for example across a cohort of students.

The fundamental building blocks of the ENA approach are the *unit of analysis* and the *conversation* (or *stanza*). The unit of analysis defines how conversations are grouped together to produce the data points for the network, while the choice of conversation determines which connections are included in a particular analysis. All concepts that appear anywhere within the same conversation are considered to be connected to one another. Typically, no weighting is used at this level – concepts are either connected or not. A conversation can be as short as an individual message, or much longer; for example, an entire discussion thread. The unit of analysis combines conversations to form the weighted network. For example, the conversations on each day might be grouped together for analysis, in order to track the development of conceptual links over time, with the unit of analysis being a day rather than a student.

We used the same network configuration to answer *both of our research questions*. We set the conversation parameter to be the message, allowing us to examine the association between framework concepts at the lowest granularity. Since we intended to look at the moderating effect of the instructional scaffolds in **RQ2**, the unit of analysis was a compound one: student within instructional condition. Students typically belonged to only one of the *Control* and *Treatment* groups but took on both the *Research Expert* and *Practising Researcher* roles.[1] By aggregating together the messages sent by a student within a single instructional condition, each student was represented in the network data twice, once for each role. We excluded from the ENA analysis the 6 messages that were sent by an instructor.

As **RQ1** aimed to better understand the general associations between the individual CoI phases of cognitive presence and the ICAP modes of cognitive engagement, we explored the mean network formed by all students taken together. To address **RQ2**, we first looked at differences between pairs of instructional conditions in a projection network. Since the same projection space was used for both the mean network and the projection network, the points in the projection network were able to be interpreted using the nodes from the mean network as 'landmarks'. We used a Mann-Whitney test to determine whether the data points from the two groups were significantly different along either of the two axes of the projection space.

We further examined how the associations between constructs differed across instructional conditions by generating additional mean networks from the relevant subsets of the data. When two different data sets are used to generate ENA networks, the latent spaces defined by their data points will differ, and it is likely that their projection spaces will have nothing in common. This means it is not possible to make direct comparisons between points in two different networks. However, a qualitative comparison is possible by observing the *relative* positioning of the nodes within each network.

5 RESULTS

5.1 Results for RQ1: general associations between the framework constructs

Our first research question asked about the general associations between individual labels in each of the frameworks; that is, between the CoI phases of cognitive presence and the ICAP modes of cognitive engagement. The distribution of labels corresponding to each of the CoI phases of cognitive presence was plotted against the labels from the extended ICAP taxonomy in Figure 2. No exclusive correspondences were found between any pairs of labels. However, a trend was visible in the chart whereby the labels indicating greater depth and quality of participation in the CoI framework appeared more frequently in conjunction with the higher ICAP indicators. For example, the number of messages with the *Integration* label increased almost linearly across the three highest ICAP modes; while *Triggering Event* messages were approximately equally likely to belong to *Active Targeted* and *Constructive Extending*.

Figure 3 shows the average ENA network for all students, using data from all four instructional conditions. The framework labels are shown using their short labels (Tables 1 and 2) to reduce visual clutter. The X axis accounts for 30.8% of the variance in the data and the Y axis accounts for 11.5%. The X axis primarily distinguishes between the interactive and non-interactive ICAP modes of cognitive engagement, while the Y axis distinguishes between the early phases (*Triggering Event* and *Exploration*) and the later phases (*Integration* and *Resolution*) of cognitive presence.

Table 4 presents the strength of each network connection from Figure 3. There were no links among pairs of codes within the same framework because each conversation (a message) had exactly one label from each framework. The strongest connections were found between the higher-order indicators across the two frameworks. The strongest of all was the link between *Interactive* mode and *Integration* phase at 0.31, followed by *Interactive* mode and *Exploration* phase at 0.22. This result is in line with the expectation that messages in both of these CoI phases of cognitive presence tend to build on the content of previous discussion contributions, as required by the definition of ICAP *Interactive* mode.

5.2 Results for RQ2: the potential moderating role of instructional scaffolds

5.2.1 Cross-tabulation. The results in Section 5.1 gave a high-level overview of the whole data set and the general associations between the framework concepts. Our second research question asked about

[1] There were a small number of exceptions. One student repeated the course and thus participated in both the *Control* and *Treatment* groups. Four students never took on the *Practising Researcher* role and two did not act as the *Research Expert*.

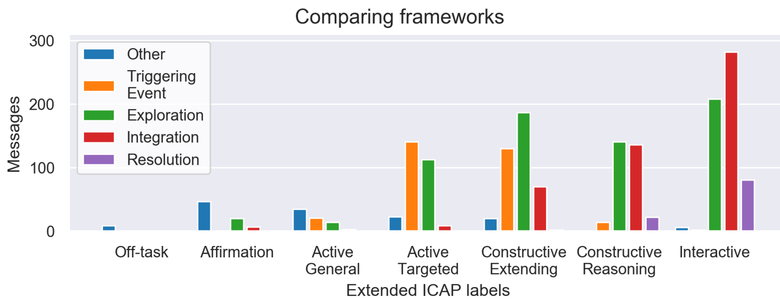

Figure 2: Message counts for each of the CoI phases of cognitive presence, broken down by the extended ICAP labels.

Table 4: ENA network weights for the network in Figure 3. Values greater than 0.10 are shown in bold.

		Other	Triggering Event	Exploration	Integration	Resolution
I	Interactive	0.00	0.00	**0.22**	**0.31**	0.09
C1	Constructive Reasoning	0.00	0.01	**0.17**	**0.15**	0.02
C2	Constructive Extending	0.02	**0.17**	**0.20**	0.09	0.00
A1	Active Targeted	0.03	**0.14**	**0.14**	0.01	0.00
A2	Active General	0.05	0.02	0.02	0.00	0.00
F	Affirmation	0.04	0.00	0.03	0.00	0.00
O	*Off-task*	0.00	0.00	0.00	0.00	0.00

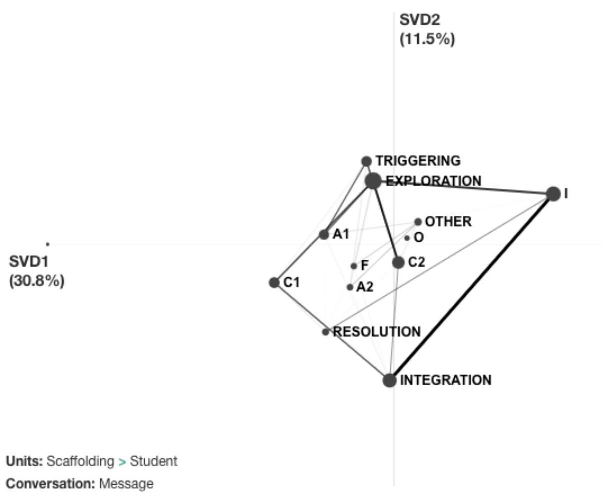

Figure 3: ENA network constructed by summing the label co-occurrences within the messages sent by each individual student in each condition and plotting the normalised distributions. The node positions indicate the locations of each framework construct in the projected space. The thickness of the lines joining the nodes indicates the average strength of the connection across all data points.

the potential moderating effects of two instructional scaffolds. To drill down into more detail we plotted heat maps for each of the four instructional conditions (Figure 4). The same general trend was seen across all four conditions as for the data set overall: the messages that received a higher indicator of depth and quality of participation on one framework were more likely to receive a higher indicator in the other framework as well, but there were no strong links between constructs at the lower levels.

One of the most striking differences between the conditions was the virtual absence of messages with the *Constructive Reasoning* label when the student was in the role of *Research Expert*, while such messages were commonly seen for students in the *Practising Researcher* role. The label definitions provide an explanation for this disparity. A message could only be labelled as *Interactive* if it was a response to the *content* of an earlier message, in addition to displaying explanation or reasoning. The first message in each thread in this data set lacked substantive content: the student simply announced the new topic and provided the link to their presentation. Thus, a message posted by a *Practising Researcher* in response to the opening message of the thread, and displaying explanation or reasoning, would be labelled as *Constructive Reasoning* rather than *Interactive*. In contrast, replies from the *Research Expert* frequently built on the content of earlier posts (*Interactive*) and only rarely got the *Constructive Reasoning* label.

5.2.2 ENA network for the whole data set. The ENA projection network in Figure 5 shows a collection of points corresponding to individual students, located based on the position of the centroid of that student's network in the projection space, and coloured according to the relevant instructional condition. The squares indicate the group means, and the dashed lines around them represent the 95% confidence intervals. The *Practising Researcher* and *Research Expert* groups were separated along the X axis, corresponding to the divide between interactive and non-interactive ICAP modes; while the *Control* and *Treatment* groups were separated along the Y axis, corresponding to the split between the earlier and later CoI phases of cognitive presence. The separation between user roles

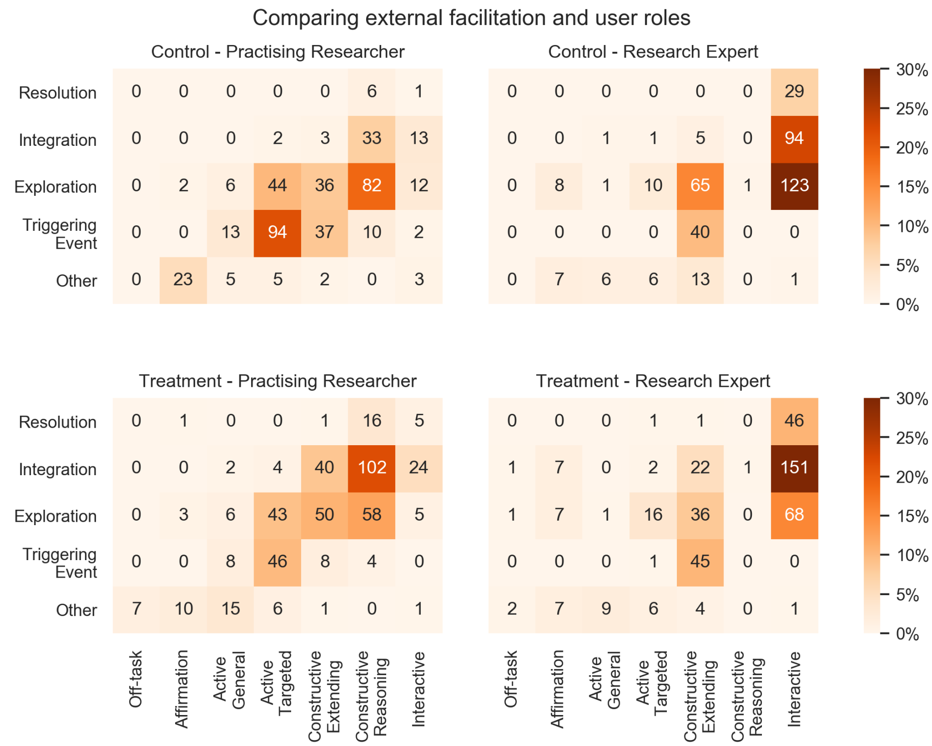

Figure 4: Distribution of labels across the two frameworks, broken down by instructional condition. The numbers in each cell indicate the raw number of messages having that combination of framework labels and the colour density indicates the distribution of labels within each condition.

was visually almost complete, with only a couple of data points from each role condition straying over the mid line. The distinction between the *Control* and *Treatment* groups appeared to be somewhat less clean. However, a series of Mann-Whitney tests showed many statistically significantly differences with large effect sizes (Table 5). In order to understand the impact of the instructional scaffolds more fully, we addressed them separately.

5.2.3 The effect of external facilitation. Separate ENA networks for the *Control* and *Treatment* groups were plotted (Figure 6) to visualise how the associations between the framework concepts varied with the external facilitation conditions. The network configuration parameters (unit of analysis and conversation) were the same as for the network built from the full data set (Figure 3), but the differences in the data resulted in the construction of different projection spaces. The amount of variance explained by the two major axes in each case was similar to the network in Figure 3 that used all the data (34.9% and 34.0% respectively for the X axis, and 12.2% and 11.8% for the Y axis). The *absolute* positions of the network nodes cannot be meaningfully compared because the networks each use a different projection space, but the *relative* positions of nodes with respect to one another within each network can be used to discover how the patterns of association between constructs differ between conditions. Specifically, the change in relative positioning of the nodes between the *Control* and *Treatment* networks indicates a shift in the associations between framework concepts due to the external facilitation intervention.

The general orientation of the networks in Figure 6 was similar to those in Figure 3, with the *Constructive Reasoning* and *Interactive* modes at opposite ends of the X axis, and the Y axis distinguishing between the earlier and later phases of cognitive presence. In the *Control* condition, the *Constructive Reasoning* mode was located near the *Triggering Event* phase; while in the *Treatment* condition it was closer to the *Resolution* phase. This suggests the existence of an association in the *Treatment* condition between using explanation and reasoning, and reaching the later phases of cognitive presence; while in the *Control* condition, despite sending such messages, students were more likely to remain in the earlier phases of cognitive presence. The *Constructive Reasoning* mode was used almost exclusively by students in the *Practising Researcher* group (Section 5.2.1), suggesting that the external facilitation intervention was successful in changing the focus of their contributions from exploring a single idea from the presentation (albeit with explanation and reasoning), to reasoning about how the presented material might relate to other content from the course.

Additionally, the *Constructive Extending* mode changed position from the right side of the network (more interactive) in the *Control* condition to the left side (less interactive) in the *Treatment* condition, while remaining nearer the early phases of cognitive presence than the later ones. One interpretation of this result is that the

Table 5: Results from a series of Mann-Whitney tests comparing conditions for statistical differences at the $\alpha = 0.05$ level along both X (SVD1) and Y (SVD2) axes, with significant results highlighted in bold. Effect sizes greater than 0.5 are considered large, while effect sizes between 0.3 and 0.5 are considered medium [6].

Axis	Condition	Comparator	*U*	*p*	*r*
SVD1	Control-Practising Researcher	Control-Research Expert	2.00	0.00	**1.00**
		Treatment-Research Expert	4.00	0.00	**1.00**
		Treatment-Practising Researcher	696.00	0.27	0.14
	Treatment-Practising Researcher	Control-Research Expert	3.00	0.00	**1.00**
		Treatment-Research Expert	6.00	0.00	**0.99**
	Control-Research Expert	Treatment-Research Expert	720.50	0.22	0.16
SVD2	Treatment-Practising Researcher	Control-Practising Researcher	353.00	0.00	**0.57**
		Control-Research Expert	386.00	0.00	**0.54**
		Treatment-Research Expert	740.00	0.04	0.25
	Treatment-Research Expert	Control-Practising Researcher	381.00	0.00	**0.54**
		Control-Research Expert	439.00	0.00	**0.49**
	Control-Research Expert	Control-Practising Researcher	657.50	0.63	0.06

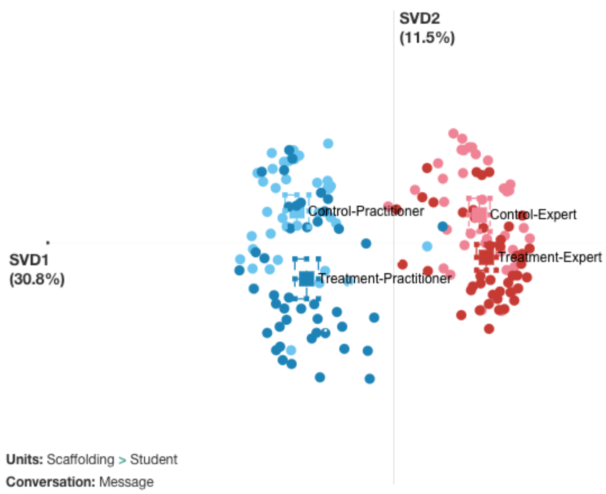

Figure 5: The ENA network constructed using a compound unit of analysis – student within instructional condition – with message as the conversation. It uses the same projection space as Figure 3 and shows the centroid of each student's network as a point. To reduce visual clutter, the *Practising Researcher* group is labelled 'Practitioner' and the *Research Expert* group is shown as 'Expert'.

students in the *Control* group, in both user roles, who tended to introduce new information to the discussion (*Constructive Extending*) were also likely to build on the contributions of others (*Interactive*); while in the *Treatment* group, the external facilitation intervention discouraged students from introducing new material without explaining how it linked to the earlier discussion – the behaviour captured by the combination of the *Constructive Extending* mode and the earlier phases of cognitive presence – and thus increased the differences between students who sent a lot of *Interactive* messages and those whose messages generally indicated lower levels of critical thinking.

5.2.4 The effect of assigning user roles. The *Practising Researcher* and *Research Expert* groups were plotted as separate ENA networks (Figure 7) to visualise how the associations between the framework concepts were moderated by the assignment of user roles. The configuration parameters were again the same as for the network constructed from the full data set (Figure 3), but the differences within the data resulted in different latent space projections – in this case leading to very different layouts compared to the earlier networks.

The network generated from the messages sent by students in the *Practising Researcher* role bore little resemblance to any of the earlier networks. The X axis accounted for 25.0% of the variance in the data and the Y axis for 18.1%. Unlike the earlier networks, the X axis primarily distinguished between the phases of cognitive presence, with the *Triggering Event* phase on the far right, the *Exploration* phase near the middle, and the *Integration* phase at the far left. The *Resolution* phase was used rarely in this subset of the data and appeared in the mid-left. The Y axis accounted for the distinction between the lower (*Active General* and *Active Targeted*) and higher (*Constructive Extending*, *Constructive Reasoning*, and *Interactive*) ICAP modes of cognitive engagement.

The *Research Expert* network shared some commonality with the *Practising Researcher* network. The phases of cognitive presence were again distributed along the X axis, accounting for 33.9% of the variance, although the positions of the *Triggering Event* and *Exploration* phases on the right of the network were reversed. The Y axis accounted for 17.0% of the variance. The *Integration* and *Exploration* phases were nearest the top, while the *Resolution* phase was at the bottom. Neither axis had a clear interpretation in terms of the ICAP modes of cognitive engagement.

We compared the relative positions of the individual constructs in the networks based on user roles in Figure 7 and the network constructed from the full data set in Figure 3. We found no clear indications of any changes to the associations between framework concepts moderated by the user role assignment intervention. In both conditions, the *Exploration* phase of cognitive presence was positioned near the ICAP *Constructive Extending* mode. In the *Practising Researcher* network, the *Constructive Reasoning* mode was also in the neighbourhood, while in the *Research Expert* network, the *Interactive* mode was nearby. This reflected the major difference in label distribution between these two conditions (Section 5.2.1) due to the requirement for an *Interactive* message to be a response

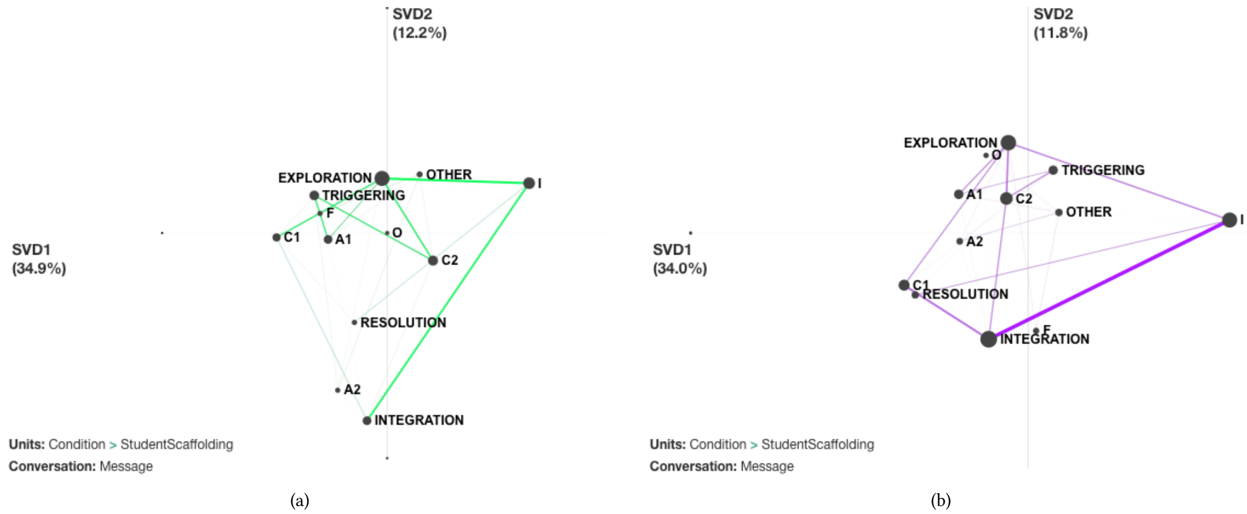

Figure 6: ENA networks constructed using only the data from (a) the *Control* condition, and (b) the *Treatment* condition.

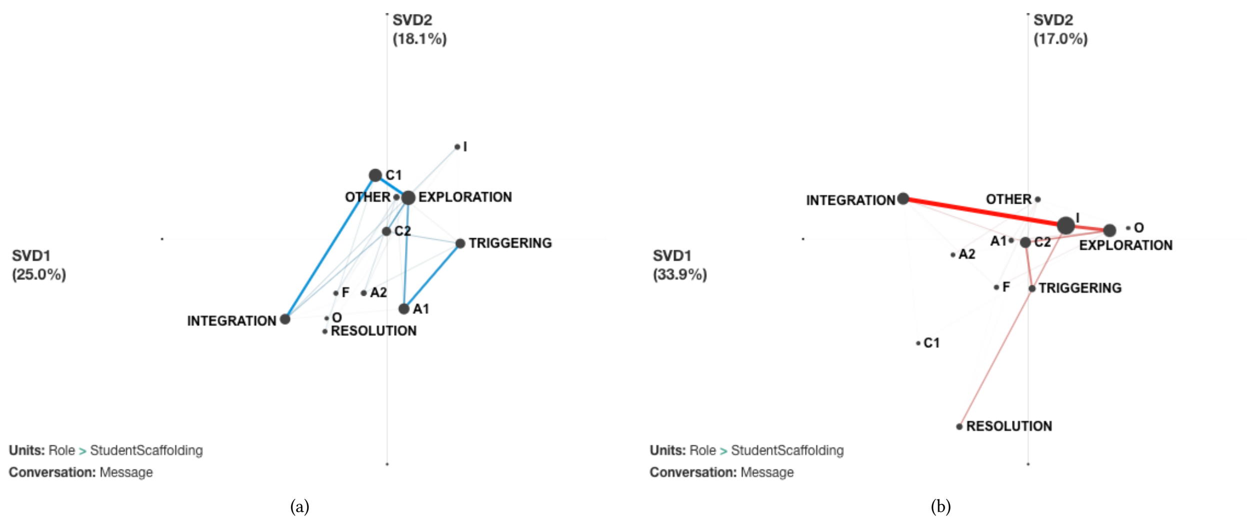

Figure 7: ENA networks constructed using only the data from (a) the *Practising Researcher* group, and (b) the *Research Expert* group.

to the substantive content of another message. Although the *Constructive Reasoning* mode in the *Practising Researcher* group was strongly linked to both the *Exploration* and *Integration* phases of cognitive presence, it was plotted in the same area of the network as *Exploration* and distant from *Integration*, indicating a greater similarity to *Exploration*. In a similar way, the *Interactive* mode in the *Research Expert* group was located near the *Exploration* phase, despite being strong linked to both *Exploration* and *Integration*. This indicates that students whose messages tended to contain explanation or reasoning were nevertheless more likely to remain in the *Exploration* phase than to progress to *Integration*. Finally, the *Constructive Extending* mode was positioned near the centre in both networks, indicating that it represented the 'conceptual average' along both axes.

6 DISCUSSION

6.1 General associations between constructs

Our analysis of the general associations between the individual CoI phases of cognitive presence and the ICAP modes of cognitive engagement in **RQ1** revealed a trend whereby the indicators of greater depth and quality of participation tended to be related while the same was not true of the indicators of lower depth and quality. This finding was evidenced both by cross-tabulation (Figure 2) and by the construction of an ENA network from the framework labels assigned to each individual message (Figure 3). This result agrees with the findings of prior work that used CoI along with Bloom's taxonomy and the SOLO taxonomy and found that the lower levels of the frameworks were "of a different nature" and thus could not be sensibly compared [28, p. 63]. Specifically, the lowest phase of cognitive presence, *Triggering Event*, indicates a message that can launch the discussion in a new direction, rather than moving towards a resolution of the current problem. In contrast, the lower modes of ICAP differ from the higher modes primarily because the contribution lacks novelty. Meanwhile, the two highest phases of cognitive presence involve bringing ideas together and constructing meaning (*Integration*) and reaching consensus (*Resolution*). The link between these and the higher modes of ICAP – *Constructive* activity to generate novel content and *Interactive* messages building on earlier contributions – is clearer.

That leaves the most common phase of cognitive presence, the *Exploration* phase, which can encompass anything from a simple paraphrase of previously shared content, through the introduction of new content, to an explanation that builds on the contribution of a partner but does not integrate multiple viewpoints. Incorporating the ICAP modes into our analysis allowed us to distinguish between these cases and thus develop a more nuanced appreciation of a student's contribution to the discussion. Future work also incorporating social presence could develop this further to look at links between the 'social climate' of the discussion [26] and the ways students build on each other's contributions.

6.2 Moderating role of instructional scaffolds

The second research question in this study, **RQ2**, aimed to explore the effects of two instructional interventions, described in Section 4.3. In each of the four conditions, the most common pairing of labels from the two frameworks was different (Figure 4).

- In the *Control-Practising Researcher* condition, where neither of the interventions was in effect, the most common label co-occurrence was the *Triggering Event* phase coupled with *Active Targeted* mode – a question relating back to information that was shared earlier in the discussion.
- In the *Control-Research Expert* condition, the most common pairing was the *Exploration* phase and *Interactive* mode – the exchange of ideas with a conversational partner – suggesting that students in the *Research Expert* role were active in driving the discussion forward.
- The most common pairing in the *Treatment-Practising Researcher* condition was the *Integration* phase with *Constructive Extending* mode, indicating the construction of meaning through explanation or reasoning. The influence of the revised assignment instructions is clear in shifting the behaviour of these students from simple clarification questions to more substantive contributions.
- Finally, in the *Treatment-Research Expert* condition, the most common co-occurrence was the *Integration* phase with *Interactive* mode, indicating the construction of meaning through explanation or reasoning in direct response to a previous message. Here, the effects of both interventions were combined, and students in the *Research Expert* role benefited from the change in behaviour of their conversational partners. This allowed the discussion to move from the *Exploration* to the *Integration* phase more frequently.

While the present study confirmed the significant shift from earlier to later phases of cognitive presence between the *Control* and *Treatment* groups in the between-subjects external facilitation intervention, noted in prior work [18, 26], there was no corresponding change in the distribution of lower to higher ICAP modes of cognitive engagement. ENA networks were used to examine and explain how the associations between the phases of cognitive presence and the ICAP modes were affected by the intervention, leading to the disparity in metrics (Figure 6).

In contrast, the role assignment intervention did not appear to have a clear moderation effect on the associations between the constructs of the two frameworks. The measures of depth and quality of participation on both frameworks were higher for the *Research Expert* group than for the *Practising Researcher* group. Further research might nevertheless reveal a moderating effect, as has been seen in work incorporating social presence [22, 26]. Future work should also consider how the association between the measures varies over time [31] and whether there is an ordering effect based on the students' experience of the different roles.

6.3 Limitations

The main limitation of this study is that it used data from only a single course, although this was collected from several offerings across an extended period of time. The relationships between the theoretical frameworks that were discovered in this study might not hold for data from a different setting; for example, the fact that participation in the forum contributed to the course grade will undoubtedly have influenced the approach taken by students.

7 CONCLUSION

The primary research contribution offered by this study is a novel approach for exploring the relationship between different indicators of the depth and quality of participation in online discussion forums. By combining insights from a combination of cross-tabulation and network analytic techniques, we uncovered connections between indicators of CoI cognitive presence and ICAP cognitive engagement and saw how student behaviour was affected by two different types of instructional intervention. The method can be applied to other situations where multiple indicators interact in potentially complex ways.

A second contribution of this study is the evaluation of how the combination of external facilitation and role assignments affected the relationship between two different measures of the depth and quality of student participation, where positive changes in one measure were not always reflected in another measure. The use of two complementary perspectives in future studies is likely to become increasingly feasible, thanks to the development of automated classifiers for both CoI [9, 11, 21, 23, 25] and ICAP [1, 20, 36]. Combining the analytic approach presented here with real-time automated labelling could allow researchers and practitioners to benefit from rich analytic insights and to evaluate interventions while a course is still in progress.

ACKNOWLEDGMENTS

This work was supported in part by the EPSRC Centre for Doctoral Training in Data Science, funded by the UK Engineering and Physical Sciences Research Council (grant EP/L016427/1) and the University of Edinburgh.

REFERENCES

[1] Thushari Atapattu, Menasha Thilakaratne, Rebecca Vivian, and Katrina Falkner. 2019. Detecting cognitive engagement using word embeddings within an online teacher professional development community. *Computers & Education* 140 (2019), 103594.

[2] John Burville Biggs and Kevin F. Collis. 1982. *Evaluating the quality of learning: the SOLO taxonomy (structure of the observed learning outcome).* Academic Press, New York.

[3] Benjamin S. Bloom, David R. Krathwohl, and Masia Bertram B. 1956. *Taxonomy of educational objectives : the classification of educational goals. Handbook 1, Cognitive domain.* Longman Group Ltd., London.

[4] Michelene T. H. Chi, Miriam Bassok, Matthew W. Lewis, Peter Reimann, and Robert Glaser. 1989. Self-Explanations: How Students Study and Use Examples in Learning to Solve Problems. *Cognitive Science* 13, 2 (1989), 145–182.

[5] Michelene T. H. Chi and Ruth Wylie. 2014. The ICAP Framework: Linking Cognitive Engagement to Active Learning Outcomes. *Educational Psychologist* 49, 4 (2014), 219–243.

[6] Jacob Cohen. 1992. A power primer. *Psychological Bulletin* 112, 1 (1992), 155–159.

[7] Bram De Wever, Tammy Schellens, Martin Valcke, and Hilde Van Keer. 2006. Content analysis schemes to analyze transcripts of online asynchronous discussion groups: A review. *Computers & Education* 46, 1 (2006), 6–28.

[8] Elaine Farrow, Johanna Moore, and Dragan Gašević. 2020. Dialogue attributes that inform depth and quality of participation in course discussion forums. In *Proceedings of the Tenth International Conference on Learning Analytics & Knowledge.* ACM, New York, NY, USA, 129–134.

[9] Elaine Farrow, Johanna Moore, and Dragan Gašević. 2019. Analysing Discussion Forum Data: A Replication Study Avoiding Data Contamination. In *Proceedings of the 9th International Conference on Learning Analytics & Knowledge.* ACM, New York, NY, USA, 170–179.

[10] Rebecca Ferguson, Zhongyu Wei, Yulan He, and Simon Buckingham Shum. 2013. An evaluation of learning analytics to identify exploratory dialogue in online discussions. In *Proceedings of the 3rd International Conference on Learning Analytics and Knowledge.* ACM Press, New York, NY, USA, 85–93.

[11] Máverick Ferreira, Vitor Rolim, Rafael Ferreira Mello, Rafael Dueire Lins, Guanliang Chen, and Dragan Gašević. 2020. Towards automatic content analysis of social presence in transcripts of online discussions. In *Proceedings of the Tenth International Conference on Learning Analytics & Knowledge.* ACM, New York, NY, USA, 141–150.

[12] Rafael Ferreira, Vitomir Kovanović, Dragan Gašević, and Vitor Rolim. 2018. Towards Combined Network and Text Analytics of Student Discourse in Online Discussions. In *Artificial Intelligence in Education*, Vol. 10947. Springer International Publishing, Cham, 111–126.

[13] Frank Fischer, Ingo Kollar, Karsten Stegmann, and Christof Wecker. 2013. Toward a script theory of guidance in computer-supported collaborative learning. *Educational Psychologist* 48, 1 (2013), 56–66.

[14] D. Randy Garrison. 2011. *E-learning in the 21st century: A framework for research and practice* (2nd ed.). Routledge, New York and London.

[15] D. Randy Garrison. 2016. *Thinking collaboratively: Learning in a community of inquiry.* Routledge, New York and London.

[16] D. Randy Garrison, Terry Anderson, and Walter Archer. 2000. Critical Inquiry in a Text-Based Environment: Computer Conferencing in Higher Education. *The Internet and Higher Education* 2, 2 (2000), 87–105.

[17] D. Randy Garrison, Terry Anderson, and Walter Archer. 2001. Critical thinking, cognitive presence, and computer conferencing in distance education. *American Journal of Distance Education* 15, 1 (2001), 7–23.

[18] Dragan Gašević, Olusola Adesope, Srećko Joksimović, and Vitomir Kovanović. 2015. Externally-facilitated regulation scaffolding and role assignment to develop cognitive presence in asynchronous online discussions. *The Internet and Higher Education* 24 (2015), 53–65.

[19] Dragan Gašević, Srećko Joksimović, Brendan R. Eagan, and David Williamson Shaffer. 2019. SENS: Network analytics to combine social and cognitive perspectives of collaborative learning. *Computers in Human Behavior* 92 (2019), 562–577.

[20] Hind Hayati, Mohammed Khalidi Idrissi, and Samir Bennani. 2020. Automatic Classification for Cognitive Engagement in Online Discussion Forums: Text Mining and Machine Learning Approach. In *Artificial Intelligence in Education*, Vol. 12164. Springer International Publishing, Cham, 114–118.

[21] Yuanyuan Hu, Claire Donald, Nasser Giacaman, and Zexuan Zhu. 2020. Towards automated analysis of cognitive presence in MOOC discussions. In *Proceedings of the Tenth International Conference on Learning Analytics & Knowledge.* ACM, New York, NY, USA, 135–140.

[22] S. Joksimović, D. Gašević, V. Kovanović, B. E. Riecke, and M. Hatala. 2015. Social presence in online discussions as a process predictor of academic performance. *Journal of Computer Assisted Learning* 31, 6 (2015), 638–654.

[23] Vitomir Kovanović, Srećko Joksimović, Zak Waters, Dragan Gašević, Kirsty Kitto, Marek Hatala, and George Siemens. 2016. Towards Automated Content Analysis of Discussion Transcripts: A Cognitive Presence Case. In *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge.* ACM, New York, NY, USA, 15–24.

[24] J. Richard Landis and Gary G Koch. 1977. The Measurement of Observer Agreement for Categorical Data. *Biometrics* 33, 1 (1977), 159–174.

[25] Valter Neto, Vitor Rolim, Rafael Ferreira, Vitomir Kovanović, Dragan Gašević, Rafael Dueire Lins, and Rodrigo Lins. 2018. Automated Analysis of Cognitive Presence in Online Discussions Written in Portuguese. In *Proceedings of the 13th European Conference on Technology Enhanced Learning*. Springer International Publishing, Cham, 245–261.

[26] Vitor Rolim, Rafael Ferreira, Rafael Dueire Lins, and Dragan Găsević. 2019. A network-based analytic approach to uncovering the relationship between social and cognitive presences in communities of inquiry. *The Internet and Higher Education* 42 (2019), 53–65.

[27] Vitor Rolim, Rafael Ferreira Leite De Mello, Vitomir Kovanovic, and Dragan Gasevic. 2019. Analysing Social Presence in Online Discussions Through Network and Text Analytics. In *2019 IEEE 19th International Conference on Advanced Learning Technologies (ICALT)*, Vol. 2161-377X. IEEE Computer Society, Los Alamitos, CA, USA, 163–167.

[28] Sarah Schrire. 2006. Knowledge building in asynchronous discussion groups: Going beyond quantitative analysis. *Computers and Education* 46, 1 (2006), 49–70.

[29] David Williamson Shaffer, Wesley Collier, and A. R. Ruis. 2016. A Tutorial on Epistemic Network Analysis: Analyzing the Structure of Connections in Cognitive, Social, and Interaction Data. *Journal of Learning Analytics* 3, 3 (2016), 9–45.

[30] Yassin Taskin, Tobias Hecking, H. Ulrich Hoppe, Vania Dimitrova, and Antonija Mitrovic. 2019. Characterizing Comment Types and Levels of Engagement in Video-Based Learning as a Basis for Adaptive Nudging. In *Proceedings of the 14th European Conference on Technology Enhanced Learning*. Springer International Publishing, Cham, 362–376.

[31] Xu Wang, Miaomiao Wen, and Carolyn P. Rosé. 2016. Towards Triggering Higher-order Thinking Behaviors in MOOCs. In *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge.* ACM, New York, NY, USA, 398–407.

[32] Xu Wang, Diyi Yang, Miaomiao Wen, Kenneth Koedinger, and Carolyn P Rosé. 2015. Investigating how student's cognitive behavior in MOOC discussion forums affect learning gains. In *Proceedings of the 8th International Conference on Educational Data Mining*. International Educational Data Mining Society, Massachusetts, USA, 226–233.

[33] Zak Waters, Vitomir Kovanović, Kirsty Kitto, and Dragan Gašević. 2015. Structure Matters: Adoption of Structured Classification Approach in the Context of Cognitive Presence Classification. In *Information Retrieval Technology*, Vol. 9460. Springer International Publishing, Cham, 227–238.

[34] Alyssa Friend Wise and Yi Cui. 2018. Unpacking the relationship between discussion forum participation and learning in MOOCs. In *Proceedings of the 8th International Conference on Learning Analytics and Knowledge.* ACM Press, New York, New York, USA, 330–339.

[35] Alyssa Friend Wise, Yi Cui, and Jovita Vytasek. 2016. Bringing order to chaos in MOOC discussion forums with content-related thread identification. In *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge.* ACM Press, New York, New York, USA, 188–197.

[36] Eran Yogev, Kobi Gal, David Karger, Marc T. Facciotti, and Michele Igo. 2018. Classifying and Visualizing Students' Cognitive Engagement in Course Readings. In *Proceedings of the Fifth Annual ACM Conference on Learning at Scale.* ACM, New York, NY, USA, 1–10.