


Technology Support for Discussion Based Learning: From Computer Supported Collaborative Learning to the Future of Massive Open Online Courses

Carolyn Penstein Rosé¹ · Oliver Ferschke¹ 

Published online: 1 March 2016

© International Artificial Intelligence in Education Society 2016

Abstract This article offers a vision for technology supported collaborative and discussion-based learning at scale. It begins with historical work in the area of tutorial dialogue systems. It traces the history of that area of the field of Artificial Intelligence in Education as it has made an impact on the field of Computer-Supported Collaborative Learning through the creation of forms of dynamic support for collaborative learning, and makes an argument for the importance of advances in the field of Language Technologies for this work. In particular, this support has been enabled by an integration of text mining and conversational agents to form a novel type of micro-script support for productive discussion processes. This research from the early part of the century has paved the way for emerging technologies that support discussion-based learning at scale in Massive Open Online Courses (MOOCs). In the next 25 years, we expect to see this early, emerging work in MOOC contexts grow into ubiquitously available social learning approaches in free online learning environments like MOOCs, or what comes next in the online learning space. These ambitious social learning approaches include Problem Based Learning, Team Project Based Learning, and Collaborative Reflection. We expect to see the capability of drawing in and effectively supporting learners of all walks of life, especially impacting currently under-served learners. To that end, we describe the current exploratory efforts to deploy technology supported collaborative and discussion-based learning in MOOCs and offer a vision for work going forward into the next decade, where we envision learning communities and open

✉ Carolyn Penstein Rosé
cprose@cs.cmu.edu

✉ Oliver Ferschke
ferschke@cs.cmu.edu

¹ Language Technologies Institute and Human-Computer Interaction Institute, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA

collaborative work communities coming together as persistent technology supported and enhanced communities of practice.

Keywords Conversational agents · Computer-supported collaborative learning · Massive open online courses

Introduction

Across theoretical frameworks in which learning is studied, including behaviorist, cognitive constructivist, and sociocultural perspectives, conversation is valued (Hmelo-Silver et al. 2013; Resnick et al. 2015). Within each of these research communities, conversation is valued in a different role. For example, it may be viewed as a stimulus for learning, a medium through which learning processes begin externally to the learner and are gradually internalized, a medium in which participation constitutes learning, or an experience where learning is development of conversational practices. A growing interest and involvement in development of technologies to support discussion for learning has been represented within the AI in Education community over the past two decades. Beginning in the mid-90s, this interest initially focused mainly on the area of tutorial dialogue systems to support individual learning (Evens and Michael 2006; Aleven et al. 2003; Zinn et al. 2002; VanLehn et al. 2002) and gradually included more and more emphasis over the past decade on computer-supported collaborative learning supported by conversational agents and other dynamic support technologies (Kumar et al. 2007; Kumar and Rosé 2011; Dyke et al. 2013; Adamson et al. 2014). With the recent rise in focus on Massive Open Online Courses (MOOCs), a new opportunity has opened for impact of technology support for collaborative and discussion-based learning experiences as innovative forms of learning in what is still mainly a solitary learning environment. In this article we will discuss both the history of development of tutorial dialogue agents in their various roles within the AI in Education community and their potential impact as one of a range of potential technological supports for discussion-based and collaborative learning in MOOCs (Ferschke et al. 2015a). Throughout we will refer to advances in the field of Language Technologies and how they have contributed to this work, and yet what is still needed from them that we envision being developed in the next 25 years. Though this paper describes a long history of work that is completed, it raises a challenge for work that will take the next 25 years to accomplish. In particular, what we would like to see is a world where learners of all walks of life, including currently under-served learners, can find ubiquitously available learning experiences to meet their personal learning goals in a way that supports them cognitively, motivationally, and socially. It is the social dimension that we focus on in this paper. Currently available platforms for learning at scale are frequently impoverished on this dimension, which is a problem for all, but especially under-served students who are in greater need of support. Learning in these environments is typically mostly a solitary experience. Though there are almost always discussion forums included in these environments, they are often just an appendage, and not effective in meeting the needs of learners, especially under-served learners who need more support.

Historical Foundations

Tutorial Dialogue

Interest in discussion-based learning in the AI in Education community is typically motivated by the famous finding that expert human tutors are highly successful at educating students (Bloom 1984; Cohen et al. 1982). Students working with an expert human tutor have been believed to achieve a learning gain of up to two standard deviations above those in a regular classroom setting. Emulating this “2 sigma effect” has long been the holy grail of intelligent tutoring research. While great strides in developing instructional technology had been made by that time especially in the area of building coached problem solving practice environments (Gertner and VanLehn 2000; Koedinger et al. 1997), achieving the goal of the full extent of the effectiveness of expert human tutors remains elusive to this day. And it should be noted that subsequent research has demonstrated that the majority of human tutors behave substantially differently from expert human tutors (Graesser et al. 1998) and that even award winning instructors do not always achieve a significantly higher learning effect than equivalent non-interactive learning support (VanLehn et al. 2007). Nevertheless, there are many impressive demonstrations of the potential of effective teacher led discussion-based learning to yield steep changes in achievement on standardized tests that persist for years and are associated with transfer effects across learning domains (Resnick et al. 2015). Thus, while identification of the most strategic way to use technology to achieve positive impact on instruction through dialogue remains something of a mystery, the desire to solve it nevertheless must be recognized to have a substantial empirical basis.

The search for the answer to the elusive “2 sigma effect” has taken many forms, but one common thread through generations of investigation has been the belief that the answer lies in the natural language dialogue that is the dominant form of communication between students and human tutors. In particular, much work has adopted a Socratic tutoring style where students are led to construct knowledge for themselves through directed questioning (Carbonell 1969; Rosé et al. 2001). Early efforts to emulate the effectiveness of human tutorial dialogue, such as the SCHOLAR system (Carbonell 1969; Carbonell 1970) and the original WHY system (Stevens and Collins 1977), were often acknowledged as the landmark systems in the history of intelligent tutoring research. Nevertheless, it was acknowledged that in that early work the conception of what Socratic tutoring is and why it should be effective was not sufficiently well worked out, and the language technologies needed to support such interactions were not yet mature.

Tutorial dialogue agents have proven capable of leading students through directed lines of reasoning in medicine (Evens and Michael 2006), Geometry (Aleven et al. 2003), electronics (Zinn et al. 2002), and physics (Freedman et al. 2000; Rosé 2000; Jordan et al. 2001; Rosé et al. 2002; Rosé et al. 2005). These tools enabled a series of successful evaluations in real classrooms, including our own work (Rosé et al. 2001; Rosé and Torrey 2005; Kumar et al. 2006). Nevertheless, it eventually became clear that one major road block to achieving impact with the technology was that student expectations of computer agents acted as a hindrance to them interacting with the agents in instructionally beneficial ways, regardless of the technical capabilities of such agents (Rosé and Torrey 2005). Thus, in recent years a growing amount of attention turned to the use of conversational agents as facilitators in collaborative learning interactions,

where the richness of human interaction could be experienced through peer interactions, and correct content and guidance could be provided by carefully designed agents (Kumar et al. 2007; Dyke et al. 2013; Adamson et al. 2014). From a different angle, some recent work still targeting individual learning with technology has sought to leverage the same feel of a group interaction through multiple agents interacting with individual students in what are referred to as dialogues (Graesser et al. 2014).

Computer-Supported Collaborative Learning

At least a decade of research shows that students can benefit from their interactions in learning groups when automated support is provided (Hmelo-Silver 2003), especially interactive and context-sensitive support (Kumar et al. 2007; Kumar and Rosé 2011; Adamson et al. 2014). Until recently, the state-of-the-art in computer supported collaborative learning has consisted of static forms of support, such as structured interfaces, prompts, and assignment of students to scripted roles (Fischer et al. 2013). This earlier work, referred to as scripted collaboration, has been a major focus of the field of Computer-Supported Collaborative Learning in the past decade, and despite its limitations, has produced numerous demonstrations of its effectiveness in improving collaborative learning.

The concept of adaptive collaborative learning support was first evaluated in a Wizard-of-Oz setup and found to be effective for supporting learning (Gweon et al. 2006). As a further proof of concept of the feasibility and potential impact of such an approach, Kumar and colleagues have evaluated tutorial dialogue agents as support for pairs working together in a power plant design task in a sophomore thermodynamics course (Kumar et al. 2007). In that study, working with a partner student and with support of an agent led to a 1.24 standard deviation improvement in learning over a control condition where students worked alone. Students who worked either just with a partner or just with the computer agent learned approximately 1 standard deviation more than the control condition students. Thus, the support offered by the partner student and the agent were synergistic rather than redundant. In other work, Cognitive Tutors provided an environment in which intelligent support scaffolded a peer tutoring process (Walker et al. 2011).

Context-sensitive or need-based support necessitates on-line monitoring of collaborative learning processes. Automatic analysis of collaborative processes is an advance in the field of Language Technologies that has value for real time assessment during collaborative learning, for dynamically triggering supportive interventions in the midst of collaborative-learning sessions, and for facilitating efficient analysis of collaborative-learning processes. The area of automatic collaborative process analysis has focused on discussion processes associated with knowledge integration. Frameworks for analysis of group knowledge building are plentiful and include examples such as Transactivity (Berkowitz and Gibbs 1983; Teasley 1997; Weinberger and Fischer 2006), Inter-subjective Meaning Making (Suthers 2006), and Productive Agency (Schwartz 1998). So far work in automated collaborative learning process analysis has focused on text-based interactions and key-click data (Soller and Lesgold 2000; Erkens and Janssen 2006; Rosé et al., 2008; McLaren et al. 2007; Mu, Stegmann, Mayfield, Rosé, & Fischer, 2012). However similar work applied to audio data has begun (Gweon et al. 2013), and in the future we anticipate automated analysis of video will enable dynamic support for learning in groups that is starting to be used in MOOC contexts, currently unsupported (Kulkarni et al. 2015).

Moocs: The Bleeding Edge

Part of the initial excitement about MOOCs was the potential to make education less expensive, and thus potentially reduce the achievement gap between low and high socio-economic status students, and ultimately improve social mobility. However, despite the fact that the majority of MOOC learners who fill out course surveys self-identify as relatively highly educated, the completion rate even among these students is an abysmally small percentage, frequently about 3 %. Analyses of attrition and learning in MOOCs even for these learners both point to the importance of social engagement for motivational support and overcoming difficulties with material and course procedures (Breslow et al. 2013). This finding echoes earlier findings that demonstrated the importance of social support for increasing retention of under-represented minorities in STEM majors in brick-and-mortar institutions (Treisman 1992). If we can improve the experience of community and social support around MOOCs, we may potentially decrease attrition for the current ilk of MOOC participants. If we can increase the available social support, our vision is that within 25 years we can create environments where under-served and under-represented learners may also find an appropriately supportive and conducive learning environment that is attractive to them and meets their instructional needs. The opportunity for impact is great if we look past the top tier and other four year institutions of higher learning, and instead target community college level instruction.¹ We must understand and respect the unique needs of such learners, however, many of which are returning to school after earlier discouraging attempts, where in many cases “life got in the way” in sensitive ways that make returning to school especially stressful. Creating a supportive environment in which these learners can find community, support, dignity, and respect is essential for achieving this impact.

However, the reality of current content-focused xMOOCs, such as the typical MOOCs offered through Coursera, edX, and Udacity, is that opportunities for exchange of ideas, help and support are limited to threaded discussion forums, which are often not well integrated with instructional activities and as a result lack many of the qualities identified as reflecting instructionally beneficial interactions from prior work in the field of CSCL (Rosé et al. 2015). In contrast, constructivist MOOCs, or cMOOCs, typically provide an eclectic variety of affordances for social interaction including blogs, Twitter communication, email, Facebook study groups and others, with the idea that students should have the freedom to find a context for learning socially within this variety that they feel comfortable with, which may be effective for engendering a wider variety of discourses contextualized within the learning and therefore meeting different instructional needs (Siemens 2005; Smith and Eng 2013). One downside of this approach, however, is that many students find the variety disorienting and anxiety-inducing, especially those who lack appropriate self-regulated learning skills (Smith and Eng 2013).

Dual-Layer Moocs: A Case Study

In Fall 2014 we launched what we termed a “dual layer” MOOC on the edX platform, with lead instructor George Siemens, and co-instructors Dragan Gašević and Ryan

¹ <http://www.clearnospace.org/blog/2014/11/18/digital-learning-research-network-dlrn/>

Baker. It was termed a “dual layer” MOOC because students had the option of following a more standard path housed in the edX platform for moving through the course in one layer or to follow a more self-directed path in an environment called ProSolo,² developed by Gašević and colleagues that formed a second layer. This MOOC, entitled Data, Analytics, and Learning (DALMOOC), was launched in October of 2014 and extended for 9 weeks in addition to an orientation week as part of in an iterative, design-based research process the author team is engaged in to integrate affordances for richer, theory motivated learning experiences into xMOOCs. The purpose of the ProSolo integration was two-fold a) increasing the social learning experiences of self-regulated learners in MOOCs through social competency approaches to learning and b) developing learner knowledge graphs that reflect what a learner knows and how they have come to know it. This ProSolo layer served as a striking alternative to the more scripted xMOOC approach to MOOC learning, which remained available to students through the edX layer. Learning in ProSolo is social in the sense that students follow and communicate with other students during the process of setting their learning objectives. However, beyond that, the learning experiences students engaged in within the ProSolo environment are still largely individual in character. Early results from DALMOOC indicate that learners experience some uncertainty and discomfort in transitioning from instructor-driven to self-regulated learning. Learners within the edX course found the distributed structure of DALMOOC confusing at times. Better scaffolding and support is needed to assist learners in transitioning between structured instruction and open engagement in social systems (either the open web or within ProSolo).

In that context, we deployed two different interventions leveraging advances in Language Technologies such as automated analysis of discussion behavior to trigger dynamic support (using an unsupervised matrix factorization approach), including what we termed the Quick Helper and a direct import of our agent supported collaborative learning technology, which we referred to as Bazaar Collaborative Reflection in this context. These serve as examples of how MOOCs can be extended to support community building and discussion based learning, just one step towards the grander vision that will take 25 years to accomplish. As just one example problem the field must seek to address going forward, the early MOOC platforms pose challenges for integration of these types of experiences. The specific examples presented here represent ways state-of-the-art learning interventions can be integrated with current MOOC platforms in a way that enables both an early form of social support as well as a platform to support further experimentation in that context. Over the next 25 years, we expect to see social learning practices supported more effectively at scale, overcoming coordination challenges, and provided ubiquitously within free learning environments, such as MOOCs.

The Quick Helper

Earlier research on MOOC discussion forums indicates that the experience of confusion as well as exposure to other students’ confusion are both associated with elevated attrition in MOOCs (Yang et al. 2015), and attempts at resolving confusion by making

² <http://dalmooc.prosolo.ca/prosolo/about.xhtml>

help requests in the threaded discussions are frequently left without a satisfactory response (Yang et al. 2014).

In response to these two problems, a specifically discussion-focused intervention, called the Quick Helper, was integrated with DALMOOC to support help seeking as well as increase the probability that help requests will be met with a satisfactory response. Results of an experimental manipulation inform how to increase willingness of learners to request helpers to participate in their threads (Howley et al. 2015). While virtually all MOOCs offer threaded discussion affordances where students can post help requests, some students are reticent to ask for help, and even when students do post help requests, many of these requests go unanswered. Our help seeking intervention connects students, whose questions may go unresolved, with student peers who may be able to answer their question. The Quick Helper is continuously available to students by means of a button. When they click, they are guided to formulate a help request. The help request is posted to the MOOC discussion board, and the text and metadata are forwarded to our Quick Helper system. Using this help request, a context-aware matrix factorization model selects three potential help providers from the pool of student peers. For this decision, it takes into account features from students, questions, and student connections as described in (Yang et al. 2014). The student is then given the option to invite one or more of these potential helpers to their thread as shown in Fig. 1. Once selected, an email with a link to the help request thread is then automatically sent to the selected helpers inviting them to participate in the thread. This approach is the first step towards bringing peer help recommendation into the MOOC space considering the unique characteristics of MOOCs such as the relatively weak social bonds between students, the great diversity among students, the lack of knowledge about specific students, and the wide range of motivations for students to participate in a MOOC which influences their possible roles as help providers. Future progress in user modeling based on the noisy MOOC data and taking into account the social ties between users within and outside of the MOOC platform will allow us move towards more sophisticated help recommendation approaches as proposed and achieved using knowledge based approaches in related work on online help recommendation for on-campus education such as i-Help (Greer et al. 1998; Vassileva et al.

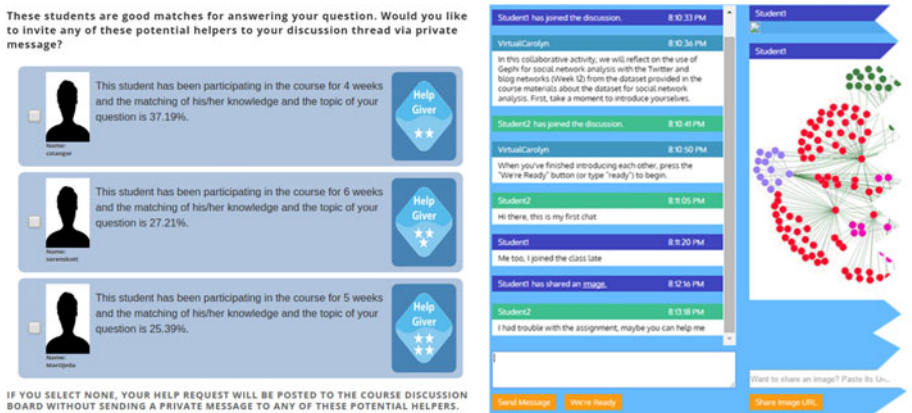


Fig. 1 A screenshot of the Quick Helper helpers selection window (left) and the Bazaar Collaborative Reflection synchronous chat (right)

2003). What QuickHelper represents is a less sophisticated form of support, but one that leverages an unsupervised matrix factorization approach to power the recommendation rather than one that requires a large scale knowledge engineering effort and extensive knowledge about learners to be modeled, in the same manner as massive recommendation affordances found in e-commerce websites.

Bazaar Collaborative Reflection

Another intervention, referred to as Bazaar Collaborative Reflection, makes synchronous collaboration opportunities available to students in a MOOC context. Research in Computer-Supported Collaborative Learning has demonstrated that conversational computer agents can serve as effective automated facilitators of synchronous collaborative learning (Dyke et al. 2013). However, typical MOOC providers do not offer students opportunities for synchronous collaboration, and therefore students have not so far been able to benefit from this technology. These activities were integrated into DALMOOC by means of a Lobby program that provides a place where students may wait briefly to be matched with a random conversational partner. Providing this service allows for some flexibility when students arrive for a chat. When students are successfully matched with a partner student within the Lobby, they are provided with a link to their own chat room. They then enter that room by clicking on the link. Then they are able to interact with each other as well as a conversational agent who appears as a regular user in the chat, as shown in Fig. 1.

While we have not yet assessed the influence of Bazaar on learning in a MOOC context, we performed a survival analysis to evaluate the influence of this collaborative learning activity on attrition. Using the number of clicks on videos and the participation in discussion forums as control variables, we found that the participation in chats lowers the risk of dropout by approximately 50 % (Ferschke et al. 2015b).

Into The Future

As we observed learners experiencing the automated and social support integrated with the innovative though fledgling DALMOOC, we observed students providing benefit to one another when they connected, both in the ProSolo Layer and the more standard edX layer. As we have noticed students gravitating towards the variety of communication media made available through ProSolo and in DALMOOC more broadly, we have begun to investigate the choices students make for where they conduct their social interaction, and what trade-offs we find in terms of the content focus and nature of discussions in the different environments. As we read their communication, we see evidence of confusion and frustration in the process of wayfinding through the choices. Thus, we find many opportunities to better support students in finding and maintaining desirable connections, support, and direction throughout these massive online courses. The use of social recommender systems (such as Quick Helper) and group collaboration tools (such as Bazaar) are expected to lead to higher levels of metacognitive monitoring, which we expect to be associated with an increase of confidence, feeling of knowing, judgment of learning, and monitoring of progress toward goals. We continue

to seek ways to integrate the technologies we have developed to achieve a more positive experience for students.

We are engaged in a design-based research process in which we grapple with these trade-offs as we seek effective practices for incorporating theory-motivated discussion-based learning opportunities in MOOCs. Specifically, we aim to import from the field of CSCL insights into the specific affordances for instructionally beneficial conversational interaction offered by alternative online discussion contexts as well as insights into what might be productive strategies for moving among them so that appropriate, integrated support and guidance for students could be designed and offered. As part of this effort, we join forces with an international collective consisting of researchers, developers, instructional designers, instructors, and policy people to forge a new vision for supporting discussion and collaboration in MOOCs.³ The goal is both to design and build platform extensions to support effective learning, but also to provide new opportunities to study learning through discussion in new ways and in novel contexts, for example, addressing important questions related to learning in highly diverse collaborative groups. Our hope is that this community will foster some of the collaborations that will produce the advances we envision in this article over the next 25 years.

Projecting the inspiration for our “dual layer MOOC” into the future, we envision MOOCs of the future as gateways into persistent Communities of Practice. Web 2.0 has produced ample evidence that with powerful architectures for supporting massive scale coordination and communication, we can engage the masses to architect resources such as Wikipedia, channel the programming skills of the many to develop open source software resources, and leverage the efforts of the everyday individual through crowdsourcing services like Amazon’s Mechanical Turk. Networked learning, where the recent phenomenon of MOOCs is one example, is a corresponding form of online participation where communities of shorter duration form for the purpose of providing a context in which individuals can better themselves through a combination of individual learning and social engagement in mentoring relationships. In all cases, the greatest resource lies in the great diversity of perspectives, expertise, skills, and energy of the masses. In the kind of MOOCs of the future we envision, communities may be launched with something like a learning progression or set of learning progressions that are more or less pre-defined like courses. However, students will learn collaboratively, through mentoring from more experienced community participants. The learning progressions will bridge learning and practice. The goal will be to engage learners in practices of Legitimate Peripheral Participation, which situate learning in some form of work (Lave and Wenger 1991). On-Demand MOOCs that are beginning to be popularized on platforms such as Coursera may contribute positively to this development by providing a more natural means for new participants to be introduced into a community dynamically over time.

Nevertheless, the great resource in crowd power of online communities comes with equally great challenges. While a diversity of perspectives and expertise offers the opportunity for achieving greater insight, richness of understanding, and heightened creativity, it may also lead to misunderstandings and conflict. Similarly, while the skills and efforts of crowds make it possible to achieve higher productivity, it may also lead to

³ DANCE: Discussion Affordances for Natural Collaborative Exchange, <http://dance.cs.cmu.edu>

diffusion of responsibility, duplicated work and even confusion. Decades of research studying the inner workings of these communities, whether focused on learning, work or a hybrid of the two, has revealed that a central problem underlying the success of these efforts is channeling the attention and efforts of the masses. This channeling may partly be achieved through facilitation of consensus and coordination achieved through a careful balance of guidance and self-direction. On the other hand, unless this is balanced with an effort to avoid premature consensus, group think, or abandonment of valuable ideas that may simply have been overlooked, support for coordination and consensus building could hinder the development of creativity or the growth of innovation.

A proposed solution to this complex and multi-faceted challenge is to provide non-coercive *personalized guidance* through awareness tools that reveal the shape and foci of the behavior data and resulting byproducts in a way that makes it natural for collaboration to grow out of individual decision-making. We have begun to lay a foundation for this work through technology that is able to identify *role-based behavior profiles* associated with productive group and community outcomes. In this lightly supervised modeling work, we do not initially prescribe the roles that should be taken up by community members. Instead through a theory guided but data driven process, we identify which configurations of behaviors are achieving positive outcomes, and then work to foster greater representation of these behavior profiles. In ongoing work in multiple learning and work communities, we will identify certain observed distributions of behaviors as contributing towards measured outcomes of interest related to productivity and learning. This modeling then will then form an effective lens for making sense of massive scale data from these contexts to fuel social recommendation agents that can aid in channeling resources where they are needed. Some initial work along these lines is just beginning (Yang et al. 2015; Ferschke et al. 2015c).

As we look to the future, we have to consider how current paradigms like script-based collaboration will adapt to the new environments for learning of the future. An important area for needed scaffolding is help exchange. Where current interventions such as the Quick Helper fall short is that they attempt to attract help resources to help requests as formulated, but it does not currently support the process of asking for help effectively, nor does it adequately and directly address the problem that certain students are reticent to ask for help. Thus, there remains much work to do in this space.

Though it is by no means the only option with current MOOC platforms, the current generation of MOOCs are mostly designed with support for transmission of knowledge, and with less attention to active engagement, and even where active engagement in learning by doing activities are featured, scant opportunities for meaningful discussion activities situated within the learning are offered, and opportunities for creation of knowledge are even more rare. Open questions remain about how to intensify the learning by doing experience and build in more support for learning through discussion and knowledge building.

For decades, the Knowledge Building community, led by Scardamalia and Bereiter, has been focused on developing scaffolding for students to engage in discussions where the goal is to create knowledge together (Scadamalia and Bereiter 2006). cMOOCs similarly aim to foster community development, with the goal of producing knowledge generation communities. However, many challenges exist: traditionally, Knowledge Forum has been mainly used in classroom settings, and there is much to learn as it is

adapted for use in a MOOC context. In Knowledge Building classrooms, the teacher plays an important role setting the culture in which Knowledge Building can take place. How can we replicate that in a MOOC? One question is related to structuring interaction. cMOOCs are loosely structured, perhaps more so than the approach to Knowledge Building that has been done in classrooms. The design space is wide open. On one end of the spectrum one might imagine a Knowledge Building MOOC, which would be “a MOOCified Knowledge Forum”. In such an environment, we might begin with the Knowledge Forum interface and pedagogy and think about how to import it into a MOOC while preserving as much as possible from current practices. On the other end of the spectrum, elements of Knowledge Building might be made available within courses structured in more typical learning progressions and employing other pedagogies as well. A third option might be to introduce the Knowledge Forum approach into an environment where large scale collaboration is already happening for producing new knowledge, such as wikis. In this paradigm, wiki activities would be embedded in MOOCs with support for knowledge building scoped within the wiki. Another issue to consider is the role of the instructor. The teacher plays a critical role in Knowledge Building classes. In order to import the associated pedagogical approach, support for the instructor role would need to be designed. Recent efforts include investigations into technology support through automated discussion analysis (Teplov et al. 2007).

The Role Of Natural Language Processing

A running theme throughout this article has been the important role of natural language processing, for making real time assessments, and for delivering interventions. Advances in the field of machine learning and language technologies, in the specific areas of text classification and text mining, as well as dialogue systems, have enabled some of the advances described above. However, the field of Language Technologies is its own field with its own history, which intersects with that of AI in Education. As we look to the future where we hope to draw more and more from this sister field, we would do well to consider its own trajectory and how it might synergize with our own going forward.

Long time members of the computational linguistics community have observed the paradigm shift that took place after the middle of the 1990s. Initially, approaches that combined symbolic and statistical methods were still of interest (Klavans and Resnick 1994). However, with the increasing focus on very large corpora and leveraging of new frameworks for large scale statistical modeling, symbolic and knowledge driven methods for natural language processing were largely left by the wayside. Along with older symbolic methods that required carefully crafted grammars and lexicons, the concept of knowledge source became strongly associated with the notion of theory, which is consistent with the philosophical notion of linguistic theory advocated by Chomskyan linguistics and other theories of formal linguistics (Steele 1990; Backofen and Smolka 1993; Wintner 2002; Schneider et al. 2004). On the positive side, the shift towards big data came with the ability to build real world systems relatively quickly. However, as knowledge-based methods were replaced with statistical models, a grounding in linguistic theory grew more and more devalued, and instead a desire to replace theory with an almost atheoretical empiricism became the zeitgeist that drove

progress within the field. As a result, work in the field focused less and less on language in ways that bridged easily to theoretical frameworks from the behavioral sciences, including the learning sciences.

About a decade ago, just as an atheoretical empiricism was becoming not only the accepted norm, but memories of a more theory driven past were sinking into obscurity, an intensive interest in analysis of social data began to grow. At first this new interest was treated as just another flavor of Information Retrieval, the area that gave us Google search. Tried and true methodologies that had worked well on retrieval of documents and management of large document collections were applied wholesale. But over time, an awareness began to grow that social data had characteristics on its own, and new methods would be needed in order to achieve adequate results. It was at this point that interest in computational social science as the answer to this need began to arise. However, the aim to replace theory with empiricism results in a culture clash when the behavioral sciences are brought into the picture, in which theory drives empiricism, and empiricism builds theory. Thus, there is no dichotomy. Atheoretical empiricism is not attractive within the behavioral sciences, including the learning sciences, where the primary value is on building theory and engaging theory in interpretation of models. In the learning sciences, process measurements must be motivated in the theory in order to be valid.

More recently, the tide has begun to turn back in ways that have the potential to benefit the learning sciences. A growing appreciation of the connection between constructs from the social sciences and work on natural language processing applied to social data is surfacing in recent years, such as data from a variety of social media environments. As this work draws significantly from research literature that informs our understanding of how social positioning takes place within conversational interactions, we are in a better position to track how students work together to create a safe or unsafe environment for knowledge building, how students take an authoritative or non-authoritative stance within a collaborative interaction, or how students form arguments in ways that either build or erode their working relationships with other students. At a very basic level, work within the area of computational sociolinguistics draws from social psychological concepts of relative power in social situations (Guinote and Vesvio 2010), in particular aspects of relative power that operate at the level of individuals in relation to specific others within groups or communities. At this level, relative power may be thought of as operating in terms of horizontal positioning, which relates to closeness and related constructs such as positive regard, trust and commitment, or vertical positioning, which relates to authority and related constructs such as approval and respect.

Within this sphere, language choices serve simultaneously as a reflection of the relative positioning of speakers to their interlocutors as well as actions that operate on those positionings (Ribeiro 2006). At a conceptual level, this work draws heavily from a foundation in linguistic pragmatics (Levinson 1983) as well as sociological theories of discourse (Gee 2011). Layered on top of the component processes for accomplishment of such goals are concepts related to expectations or norms that provide the foundation for claiming such positions, which may be described from a philosophical perspective or a sociological one (Postemes et al. 2000). Once such norms are set in place, actions performed by utterances can be judged as consistent with these norms or not. In so doing, it relates both to aspects of display of social and ethnic identity in so

far as these constructed or indexed identities are associated with rights and obligations within interaction (Myers-Scotton 1993). Here the goal is not to reflect a social or ethnic identity for its own sake, but rather to leverage those identities in order to accomplish something related to social positioning. In particular, it may position content within a configuration of sides, in terms of its status within a debate as consistent with one side or the other, or under consideration or not (Martin and White 2005), which may include the reflection of the attitude of an individual towards some entity (Martin and White 2005), or it may go beyond this reflection in attempting to persuade another and in so doing to operate on this positioning. It may also touch upon more practical goals to accomplish moves such as requesting an action or working to persuade someone of some idea. In viewing interaction as providing a context in which information and action may flow towards accomplishment of such goals, speakers position themselves and others as sources or recipients of such information and action (Martin and Rose 2007). When performatives break norms related to social position, they have implications for relational constructs such as politeness (Brown and Levinson 1987), which codifies rhetorical strategies for acknowledging and managing relational expectations while seeking to accomplish extra-relational goals.

While this very general framing provides an umbrella under which work within this area can be placed, the specific framing of individual publications varies substantially. At one end of the spectrum, some contemporary publications in the Language Technologies field are motivated purely from within the Language Technologies literature with a focus either on categories of conversational behavior defined in earlier research (Zhai and Williams 2014; Hasan and Ng 2014; Nguyen et al. 2012), or framed from purely a task related perspective (Bhatia et al. 2014) or a modeling perspective (Paul 2012). Other work has its foundations in sociolinguistic analyses of dialogue structure (Prabhakaran et al. 2012; Mayfield and Rosé 2011), Rhetorical Structure Theory (Allen et al. 2014), the Pragmatics literature on Speech-Act theory (Cadilhac et al. 2013; Ferschke et al. 2012) or Politeness theory (Bramsen et al. 2011), and Critical Discourse Analysis (Mukherjee et al. 2013). In recent years, a substantial amount of work in this area has leveraged or at least references theoretical frameworks from Social Psychology (Bak et al. 2014; Bracewell et al. 2012). Within that sphere, some work has focused on more specific subcommunities such as the area of doctor-patient interaction (Wallace et al. 2013) or the literature on social support and health (Mayfield et al. 2013).

As this bridging between fields continues to grow, we are in a better position to use theoretical insights into the deep inner workings of language to motivate design decisions that underlie the computational models we build. These insights have the potential to yield greater validity in the work on learning analytics applied to conversational interaction data. Through more intensive collaboration between communities, researchers in the Learning Sciences would have the opportunity to challenge the field of Language Technologies with pointers to constructs that would be valuable to model and predict, at the same time offering additional insight into the underlying social processes the language patterns reflect. If the conversation is successful, the field of Language Technologies will better fill its own often stated goal to provide modeling tools in service of behavioral research, including learning research. With the rise of the MOOCs, more interest in big data in education has taken root in the Language Technologies and related Data Mining communities. However, the exchange between

communities remains at a fledgling stage. In 25 years we hope to see this exchange come to fruition.

Reflections

In late 2012, the New York Times⁴ declared 2013 “the Year of the MOOC.” Massive Open Online Courses were said to overturn the century-old model of higher education and deliver top-tier teaching to the masses around the world – for free. Though this view has fallen out of favour, the most enthusiastic advocates believed that in 50 years, 10 institutions could be responsible for delivering higher education.⁵ As the year neared its end, all the excitement seemed to have vanished and the media coverage turned from ecstatic to depressed and disillusioned. Even though the hype seems to be over, the future of MOOCs has just begun now that fields such as Artificial Intelligence in Education and Computer-Supported Collaborative Learning are just starting to extend their impact to this new realm of massive online education. In this article we have set a vision for the future of technology supported collaborative and discussion-based learning at scale.

A major reason why MOOCs did not meet the high expectations is the fact that early MOOCs were mere translations of their teacher-centered face-to-face counterparts into an electronic format that ignored the importance of personal interaction in learning – both between faculty and students and among students. MOOCs of the future should not regard their students as anonymous masses of isolated individual learners who are provided with static learning materials in the form of videos and scripts but rather as communities of practice that develop their own means of support and knowledge exchange within a framework and with methods provided by the MOOC. Such a framework will include methods for assisting help exchange among students as well as means to navigate the massive amount of data that is produced by the community. Dynamic and user-centered support based on role-based behavior profiles can build on existing well-established approaches to scripted learning interaction and make use of current developments in Language Technologies to provide support on a massive scale.

Rather than regarding MOOCs as scaled-up virtual classrooms, they should be compared to interactive textbooks around which communities of practice are formed that engage in collaborative learning and cooperative work driven by a common interest. The MOOC platform plays the role of providing affordances that enable learning and community building. These MOOC platforms will continue to be seeded with learning materials and learning tasks by the MOOC designers, but as students take a more active role as MOOCs grow into full fledged communities, the shape the learning takes will be molded more and more by the community and its contributions over time. Only then do we have a chance to witness a development similar to the success story of Wikipedia, which started as a small expert-driven community and could only grow to the size it is today by opening up to the public, providing

⁴ <http://www.nytimes.com/2012/11/04/education/edlife/massive-open-online-courses-are-multiplying-at-a-rapid-pace.html>

⁵ http://www.wired.com/2012/03/ff_aiclass/

affordances for cooperative work and thereby inviting users from all over the world to build a community around a shared goal.

The current generation of MOOCs have sometimes been referred to as the Alta Vista of MOOCs, and we have not yet seen the Google Search of MOOCs. By the end of the next 25 years, we expect that the very idea of a MOOC per se will have faded into a distant memory, and will be replaced with a new, far more effective paradigm for learning at scale. Our vision is that this paradigm will be one where ambitious social learning practices will be ubiquitously offered, such as Problem Based Learning, Team Based Learning, Collaborative Reflection, and spontaneous personalized mentoring. These new environments will be active, thriving, communities where students are supported on their personal learning path.

Acknowledgments This work was funded in part through NSF grants OMA-0836012 and IIS-1320064 as well as funding from the Gates Foundation Digital Learning Research Network and funding from Google and Bosch. The authors gratefully acknowledge useful discussions with collaborators George Siemens, Dragan Gašević, Ryan Baker, and Candace Thille.

References

- Adamson, D., Dyke, G., Jang, H. J., & Rosé, C. P. (2014). Towards an agile approach to adapting dynamic collaboration support to student needs. *International Journal of AI in Education*, 24(1), 91–121.
- Aleven, V., Koedinger, K. R., & Popescu, O. (2003). A Tutorial Dialogue System to Support Self-Explanation: Evaluation and Open Questions. In U. Hoppe, F. Verdejo, & J. Kay (Eds.), *Proceedings of the 11th International Conference on Artificial Intelligence in Education, AI-ED 2003* (pp. 39–46). Amsterdam: IOS.
- Allen, K., Carenini, G., Ng, R. (2014). Detecting Disagreement in Conversations using Pseudo-Monologic Rhetorical Structure. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1169–1180 Doha, Qatar.
- Backofen, R. & Smolka, G. (1993). A complete and recursive feature theory, *Proceedings of the 31st Annual Meeting of the Association for Computational Linguistics*, pp 193–200.
- Bak, J., Lin, C., & Oh, A. (2014). Self-disclosure topic model for classifying and analyzing Twitter conversations. In *Proceedings of the International Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*, Doha, Qatar (p. 2014).
- Berkowitz, M. W., & Gibbs, J. C. (1983). Measuring the developmental features of moral discussion. *Merrill-Palmer Quarterly*, 29(4), 399–410.
- Bhatia, S., Biyani, P, & Mitra, P. (2014). Summarizing online forum discussions – can dialog acts of individual messages help? *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2127–2131.
- Bloom, B. S. (1984). The 2 Sigma Problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*, 13, 4–16.
- Bracewell, D. B., Tomlinson, M., Wang, H. (2012). A motif approach for identifying pursuits of power in social discourse. In *Proceedings of the Sixth International Conference on Semantic Computing (ICSC)* (pp. 1–8).
- Bramsen, P., Escobar-Molano, M., Patel, A. & Alonso, R. (2011). Extracting social power relationships from natural language. *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, pages 773–782.
- Breslow, L., Pritchard, D., de Boer, J., Stump, G., Ho, A., & Seaton, D. (2013). Studying learning in the worldwide classroom: Research into edX's first MOOC. *Research & Practice in Assessment*, 8, 13–25.
- Brown, P., & Levinson, S. (1987). *Politeness: some universals in language usage*. Cambridge: Cambridge University Press.

- Cadilhac, A., Asher, N., Benamara, F., & Lascarides, A. (2013). Grounding strategic conversation: using negotiation dialogues to predict trades in a win-lose game. In *Proceedings of Empirical Methods in Natural Language Processing (EMNLP)* (pp. 357–368).
- Carbonell, J. (1969). On man-computer interaction: a model and some related issues. *IEEE Transactions on Systems Science and Cybernetics*, 5(1), 16–26.
- Carbonell, J. (1970). AI in CAI: an artificial intelligence approach to computer-assisted instruction. *IEEE Transactions on Man-Machine Systems*, 11(4), 190–202.
- Cohen, P. A., Kulik, J. A., & Kulik, C. C. (1982). Educational Outcomes of Tutoring: A meta-analysis of Findings. *American Educational Research Journal*, 19, 237–248.
- Dyke, G., Adamson, A., Howley, I., & Rosé, C. P. (2013). Enhancing scientific reasoning and discussion with conversational agents. *IEEE Transactions on Learning Technologies*, 6(3), 240–247. Special issue on Science Teaching.
- Erkens, G., & Janssen, J. (2006). Automatic coding of communication in collaboration protocols. Proceedings of the 7th International Conference of the Learning Sciences (ICLS).
- Evens, M., & Michael, J. (2006). *One-on-One tutoring by humans and machines*. Lawrence Erlbaum and Associates: Mahwah.
- Ferschke, O., Gurevych, I., & Chebotar, Y. (2012). Behind the Article: Recognizing Dialog Acts in Wikipedia Talk Pages, in *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pp 777–786
- Ferschke, O., Howley, I., Tomar, G., Yang, D., Rosé, C. P. (2015a). Fostering Discussion across Communication Media in Massive Open Online Courses, *Proceedings of Computer Supported Collaborative Learning*. Pages 459–466.
- Ferschke, O., Yang, D., Tomar, G., Rosé, C. P. (2015b). Positive Impact of Collaborative Chat Participation in an edX MOOC, *Proceedings of AI in Education*. Pages 115–124.
- Ferschke, O., Yang, D., Rosé, C. P. (2015c). A Lightly Supervised Approach to Role Identification in Wikipedia Talk Page Discussions. *Proceedings of the International AAAI Conference on Weblogs and Social Media. Workshop on Wikipedia, a Social Pedia: Research Challenges and Opportunities*. Pages 43–47.
- Fischer, F., Kollar, I., Stegmann, K., & Wecker, C. (2013). Toward a script theory of guidance in computer-supported collaborative learning. *Educational Psychologist*, 48(1), 56–66.
- Freedman, R. K., Rosé, C. P., Ringenberg, M. A., VanLehn, K. (2000). ITS Tools for Natural Language Dialogue: A Domain Independent Parser and Planner, *Proceedings of the 5th International Conference on Intelligent Tutoring Systems*, pp 433–442.
- Gee, J. P. (2011). *An Introduction to Discourse Analysis: Theory and Method, Third Edition*. New York: Routledge.
- Gertner, A., & VanLehn, K. (2000). Andes: A Coached Problem Solving Environment for Physics. In G. Gauthier, C. Frasson, & K. VanLehn (Eds.), *Intelligent Tutoring Systems: 5th International Conference. Lecture Notes in Computer Science, Vol. 1839* (pp. 133–142). Berlin: Springer.
- Graesser, A. C., Bowers, C. A., Hacker, D. J., & Person, N. K. (1997). An anatomy of naturalistic tutoring. In K. Hogan & M. Pressley (Eds.), *Scaffolding student learning: Instructional approaches and issues*. Cambridge: Brookline Books.
- Graesser, A., Li, H., Forsyth, C. (2014). Learning by Communicating in Natural Language with Conversational Agents. *Current Directions in Psychological Science* 23(5):374–380
- Greer, J., McCalla, G., Cooke, J., Collins, J., Kumar, V., Bishop, A., & Vassileva, J. (1998). The Intelligent HelpDesk: Supporting Peer Help in a University Course. In *Proceedings of the 4th International Conference on Intelligent Tutoring Systems, San Antonio, TX, LNCS No.1452* (pp. 494–503). Berlin: Springer Verlag.
- Guinote, A., & Vesvio, T. (2010). *The social psychology of power*. The Guilford Press: New York.
- Gweon, G., Rosé, C. P., Zais, Z., & Carey, R. (2006). Providing Support for Adaptive Scripting in an On-Line Collaborative Learning Environment. In *Proceedings of CHI 06: ACM conference on human factors in computer systems* (pp. 251–260). New York: ACM Press.
- Gweon, G., Jain, M., Mc Donough, J., Raj, B., & Rosé, C. P. (2013). Measuring prevalence of other-oriented transactive contributions using an automated measure of speech style accommodation. *International Journal of Computer Supported Collaborative Learning*, 8(2), 245–265.
- Hasan, K. & Ng, V. (2014). Why are You Taking this Stance? Identifying and Classifying Reasons in Ideological Debates. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 751–762.
- Hmelo-Silver, C. E. (2003). Analyzing collaborative knowledge construction: multiple methods for integrated understanding. *Computers & Education*, 41(4), 397–420.

- Hmelo-Silver, C., Chinn, C., Chan, C., & O'Donnell, A. (2013). *The international handbook of collaborative learning*. Routledge: New York.
- Howley, I., Tomar, G., Yang, D., Ferschke, O., Rosé, C.P. (2015). Alleviating the Negative Effect of Up and Downvoting on Help Seeking in MOOC Discussion Forums. *Proceedings of the 17th International Conference on Artificial Intelligence in Education*. pp. 629–632.
- Jordan, P., Rosé, C. P., & Vanlehn, K. (2001). *Tools for Authoring Tutorial Dialogue Knowledge, Proceedings of the 10th International Conference on AI in Education*. Texas: San Antonio.
- Klavans, J., & Resnick, P. (1994). *The balancing act: combining symbolic and statistical approaches to language*. The MIT Press: Cambridge.
- Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. A. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education*, 8, 30–43.
- Kulkarni, C., Cambre, J., Kotturi, Y., Bernstein, M.S., & Klemmer, S.R. (2015). Talkabout: Making Distance Matter with Small Groups in Massive Classes. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*. ACM: New York (pp. 1116–1128).
- Kumar, R., & Rosé, C. P. (2011). Architecture for building conversational agents that support collaborative learning. *IEEE Transactions on Learning Technologies*, 4(1), 21–34.
- Kumar, R., Rosé, C. P., Aleven, V., Iglesias, A., & Robinson, A. (2006). *Evaluating the Effectiveness of Tutorial Dialogue Instruction in an Exploratory Learning Context, ITS'06 Proceedings of the 8th international conference on Intelligent Tutoring Systems*, pp 666–674. Berlin, Heidelberg: Springer-Verlag.
- Kumar, R., Rosé, C. P., Wang, Y. C., Joshi, M., Robinson, A. (2007). Tutorial Dialogue as Adaptive Collaborative Learning Support, *Proceedings of the 2007 conference on Artificial Intelligence in Education: Building Technology Rich Learning Contexts That Work*, pp 383–390
- Lave, J., & Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*. Cambridge: Cambridge University Press.
- Levinson, S. C. (1983). *Pragmatics. Cambridge Textbook in Linguistics*. Cambridge: Cambridge University Press.
- Martin, J. R., & Rose, D. (2007). *Working with discourse: Meaning beyond the clause*. London: Continuum.
- Martin, J. R., & White, P. R. R. (2005). *The language of evaluation: Appraisal in English*. Basingstoke: Palgrave Macmillan.
- Mayfield, E. & Rosé, C. P. (2011). Recognizing Authority in Dialogue with an Integer Linear Programming Constrained Model. *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, pages 1018–1026
- Mayfield, E., Adamson, D., & Rosé, C. P. (2013). Recognizing Rare Social Phenomena in Conversation: Empowerment Detection in Support Group Chatrooms. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics* (pp. 104–113).
- McLaren, B., Scheuer, O., De Laat, M., Hever, R., de Groot, R. & Rosé, C. P. (2007). Using Machine Learning Techniques to Analyze and Support Mediation of Student EDiscussions. *Proceedings of the 2007 conference on Artificial Intelligence in Education: Building Technology Rich Learning Contexts That Work* (pp. 331–338).
- Mu, J., Stegmann, K., Mayfield, E., Rosé, C. P., & Fischer, F. (2012). The ACODEA Framework: Developing Segmentation and Classification Schemes for Fully Automatic Analysis of Online Discussions. *International Journal of Computer Supported Collaborative Learning* 7(2), pp285-305.
- Mukherjee, A., Venkataraman, V., Liu, B., & Meraz, S. (2013). Public Dialogue: Analysis of Tolerance in Online Discussions. *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, pages 1680–1690.
- Myers-Scotton, C. (1993). *Social motivations for codeswitching: Evidence from Africa. Oxford studies in language contact*. Oxford: Clarendon Press.
- Nguyen, V., Boyd-Graber, J., & Resnick, P. (2012). SITS: A Hierarchical Nonparametric Model using Speaker Identity for Topic Segmentation in Multiparty Conversations. *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics*, pages 78–87.
- Paul, M. (2012). Mixed Membership Markov Models for Unsupervised Conversation Modeling. *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 94–104, Jeju Island, Korea, 12–14 July 2012.
- Postemes, T., Spears, R., & Lea, M. (2000). The Formation of Group Norms in Computer-Mediated Communication. *Human Communication Research*, 26(3), 341–371.
- Prabhakaran, V., Rambow, O., Diab, M. (2012). Predicting Overt Display of Power in Written Dialogs. *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 518–522.

- Resnick, L., Asterhan, C., & Clarke, S. (2015). *Socializing Intelligence through Academic Discourse*. Washington, DC: American Educational Research Association.
- Ribeiro, B. (2006). Footing, positioning, voice: Are we talking about the same thing? In A. Fina, D. Schiffrin, & M. Bamberg (Eds.), *Discourse and Identity* (pp. 48–82). Cambridge: Cambridge University Press.
- Rosé, C. P. (2000). A Framework for Robust Semantic Interpretation, *Proceedings of 1st Meeting of the North American Chapter of the Association for Computational Linguistics*, pp 311–318
- Rosé, C. P., & Torrey, C. (2005). Interactivity versus Expectation: Eliciting Learning Oriented Behavior with Tutorial Dialogue Systems, *Proceedings of 10th IFIP TC13 International Conference on Human-Computer Interaction - Interact 2005*, Lecture Notes in Computer Science Volume 3585, pp 323–336.
- Rosé, C. P., Moore, J. D., VanLehn, K., & Allbritton, D. (2001). *A Comparative Evaluation of Socratic versus Didactic Tutoring*, *Proceedings of the 23rd Annual Conference of the Cognitive Sciences Society*, pp 869–874. Scotland, UK: Edinburgh.
- Rosé, C. P., Roque, A., Bhembe, D., & VanLehn, K. (2002). *An Efficient Incremental Architecture for Robust Interpretation*, *HLT '02 Proceedings of the 2nd International Conference on Human Languages Technologies*, pp 307–312. California: San Diego.
- Rosé, C. P., Alevan, V., Carey, R., Robinson, A., Wu, C. (2005). A First Evaluation of the Instructional Value of Negotiatble Problem Solving Goals on the Exploratory Learning Continuum, *Proceedings of the 2005 conference on Artificial Intelligence in Education: Supporting Learning through Intelligent and Socially Informed Technology*, pp 563–570.
- Rosé, C. P., Goldman, P., Sherer, J. Z., Resnick, L. (2015). Supportive Technologies for Group Discussion in MOOCs, *Current Issues in Emerging eLearning*, Special issue on MOOCs, January 2015.
- Rosé, C. P., Wang, Y. C., Cui, Y., Arguello, J., Stegmann, K., Weinberger, A., & Fischer, F. (2008). Analyzing Collaborative Learning Processes Automatically: Exploiting the Advances of Computational Linguistics in Computer-Supported Collaborative Learning, submitted to the *International Journal of Computer Supported Collaborative Learning* 3(3), pp237-271.
- Scadamalia, M., & Bereiter, C. (2006). Knowledge Building: Theory, Pedagogy, and Technology. In K. Sawyer (Ed.), *Cambridge Handbook of the Learning Sciences* (pp. 97–118). New York: Cambridge University Press.
- Schneider, G., Dowdall, J., & Rinaldi, F. (2004). A robust and hybrid deep-linguistic theory applied to large-scale parsing, *Proceedings of the 3rd Workshop on Robust Methods in Analysis of Natural Language Data*, pp 14–23.
- Schwartz, D. (1998). The productive agency that drives collaborative learning. In P. Dillenbourg (Ed.), *Collaborative learning: Cognitive and computational approaches* (pp. 197–218). New York: Elsevier.
- Siemens, G. (2005). Connectivism: a learning theory for a digital age. *International Journal of Instructional Technology and Distance Learning*, 2(1), 3–10.
- Smith, B. & Eng, M. (2013). MOOCs: A Learning Journey: Two continuing education practitioners investigate and compare cMOOC and xMOOC learning models and experience, in *Proceedings of the 6th International Conference on Hybrid Learning and Continuing Education*, pp. 244–255.
- Soller, A., & Lesgold, A. (2000). Modeling the Process of Collaborative Learning. *Proceedings of the International Workshop on New Technologies in Collaborative Learning*.
- Steele, J. (1990). Meaning-Text Theory: Linguistics, Lexicography, and Implications. *Computational Linguistics*, 18(1), 108–110.
- Stevens, A., & Collins, A. (1977). The Goal Structure of a Socratic Tutor. In *Proceedings of the National ACM Conference*. Association for Computing Machinery, New York, (Also available as *BBN Report No. 3518* from Bolt Beranek and Newman Inc., Cambridge, Mass., 02138).
- Suthers, D. (2006). Technology affordances for inter-subjective meaning making: a research agenda for CSCLE. *International Journal of Computer Supported Collaborative Learning*, 1, 315–337.
- Teasley, S. D. (1997). Talking about reasoning: How important is the peer in peer collaboration? In L. B. Resnick, R. Säljö, C. Pontecorvo, & B. Burge (Eds.), *Discourse, tools and reasoning: Essays on situated cognition* (pp. 361–384). Berlin: Springer-Verlag.
- Teplövs, C., Donohue, Z., Scadamalia, M., & Philip, D. (2007). Tools for concurrent, embedded, and transformative assessment of knowledge building processes and progress. In *Proceedings of the 8th international conference on Computer supported collaborative learning (CSCL'07)*. C. A. Chinn, G. Erkens, S. Puntambekar (Eds.). *International Society of the Learning Sciences* (pp. 721–723).
- Treisman, U. (1992). Studying Students Studying Calculus: A Look at the Lives of Minority Mathematics Students in College. *The College Mathematics Journal*, 23(5), 362–372.
- VanLehn, K., Jordan, P., Rosé, C. P., & The Natural Language Tutoring Group (2002). The Architecture of Why2-Atlas: a coach for qualitative physics essay writing. In S. A. Cerri, G. Gouarderes, & F. Paraguacu (Eds.), *Intelligent Tutoring Systems, 2002: 6th International Conference* (pp. 158–167). Berlin: Springer.

- VanLehn, K., Graesser, A., Jackson, G. T., Jordan, P., Olney, A., & Rosé, C. P. (2007). Natural Language Tutoring: A comparison of human tutors, computer tutors, and text. *Cognitive Science*, *31*(1), 3–52.
- Vassileva, J., McCalla, G., Greer, J. (2003). Multi-Agent Multi-User Modeling in I-Help. *User Modeling and User Adapted Interaction* *13*(1), pp.179–210. Kluwer Academic Publishers.
- Walker, E., Rummel, N., & Koedinger, K. R. (2011). Designing automated adaptive support to improve student helping behaviors in a peer tutoring activity. *International Journal of Computer-Supported Collaborative Learning*, *6*, 279–306.
- Wallace, B.C., Trikalinos, T. A., Laws, M. B., Wilson, I. B., & Charniak, E. (2013). A generative joint, additive, sequential model of topics and speech acts in patient-doctor communication. In EMNLP 2013–2013 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference (pp. 1765–1775). Association for Computational Linguistics (ACL).
- Weinberger, A., & Fischer, F. (2006). A framework to analyze argumentative knowledge construction in computer-supported collaborative learning. *Computers & Education*, *46*(1), 71–95.
- Wintner, S. (2002). Formal language theory for natural language processing, in *Proceedings of the Workshop on Effective Tools and Methodologies for Teaching Natural Language Processing and Computational Linguistics*, pp 71–76.
- Yang, D., Piergallini, M., Howley, I., Rose, C. (2014). Forum thread recommendation for massive open online courses. *Proceedings of 7th Intl Conf. on Educational Data Mining*. pp 257–260.
- Yang, D., Wen, M., Rosé, C. P. (2015). Weakly Supervised Role Identification in Teamwork Interactions. *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics*. Pages 1671–1680.
- Zhai, K., & Williams, J. (2014). Discovering latent structure in task-oriented dialogues. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, pp. 36–46. Baltimore, Maryland, USA.
- Zinn, C., Moore, J. D., & Core, M. G. (2002). A 3-Tier Planning Architecture for Managing Tutorial Dialogue. In S. A. Cerri, G. Gouardères, & F. Paraguaçu (Eds.), *Proceedings of the Sixth International Conference on Intelligent Tutoring Systems, ITS 2002* (pp. 574–584). Berlin: Springer Verlag.