

Deconstructing the Discussion Forum: Student Questions and Computer Science Learning

Mickey Vellukunnel¹
Jeffrey Forbes³

Philip Buffum²
Sarah Heckman²

Kristy Elizabeth Boyer¹
Ketan Mayer-Patel⁴

¹Department of Computer & Information Science & Engineering, University of Florida, Gainesville, Florida

²Department of Computer Science, North Carolina State University, Raleigh, North Carolina

³Department of Computer Science, Duke University, Durham, North Carolina

⁴Department of Computer Science, University of North Carolina, Chapel Hill, North Carolina

¹{m.vellukunnel,keboyer}@ufl.edu ²{psbuffum,sarah_heckman}@ncsu.edu

³forbes@cs.duke.edu ⁴kmp@cs.unc.edu

ABSTRACT

Online discussion forums are widely used and hold great promise for supporting students in learning computer science. Understanding how we can best support students in learning computer science through online discussion forums is an important open question for the CS Ed community. This paper analyzes discussion forum posts from 395 students enrolled in CS2 across two different universities. The results demonstrate that students use the discussion forums often for logistical and relatively shallow questions. However, the largest portion of questions reflect some level of constructive problem-solving activity, and are positively correlated with course grades. Questions that neither describe students' reasoning nor their attempts to solve the problem constitute the smallest percentage of questions, but these questions may be particularly important to attend to because of their relationship to students' prior experience.

Keywords

Computer science education research; CS2 discussion forums; Piazza; dialogue analysis; question asking

1. INTRODUCTION

Online discussion forums are becoming increasingly common in computer science education because they provide an opportunity for students to learn outside of the classroom on their own time. Discussion forums provide an alternative to asking questions in person where students may feel lower pressure [19] and encounter fewer social barriers [12]. These concerns, and related emotional factors such as fear [13] and the need for affirmation [15] play important roles in students' motivation for asking—or not asking—questions. The lower pressure facilitated by discussion forums may contribute to improved outcomes in CS [18].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGCSE '17, March 08-11, 2017, Seattle, WA, USA

© 2017 ACM. ISBN 978-1-4503-4698-6/17/03...\$15.00

DOI: <http://dx.doi.org/10.1145/3017680.3017745>

One concern often voiced about discussion forums is that only a small percentage of students actively engage in them. Even when the courses require the students to post on discussion forums as part of course credit, there is evidence that the students post only the minimal amount to receive credit [10]. Research on MOOCs shows that a very small percentage of students are responsible for the vast majority of posts on the discussion forum; however, a much larger proportion of students view the forum [3].

Interestingly, students in computer science classes may participate in discussion forums more actively than prior findings might suggest. We have examined the (voluntary) online discussion forum engagement of students in CS2 across two different large universities in the southeast US, all of which used the Piazza online discussion forum which supports questions, answers, and discussions.

The data show that more than 80% of the 560 students enrolled in these classes actively posted to the discussion forum at least once via question, answer, or note (not including other forum contributions such as follow-up discussions or comments, poll answers, or edits). This high participation rate may be influenced by the teaching staff, who encouraged students to use Piazza. Additionally, the project-heavy curriculum offered many opportunities for students to encounter obstacles. Under these conditions, the online discussion forum emerged as a highly active venue where students sought help.

Because of discussion forums' potential role in supporting student learning, it is important to understand the ways in which students are using this resource. This paper makes three contributions. First, it analyzes a cross-institutional data set from CS2 students with regard to their discussion forum usage, course performance, and prior experience in computer science. Second, it presents an analysis of online questions that leverages techniques from dialogue analysis and active learning research. Finally, it demonstrates that students in CS2 use discussion forums often for logistical and relatively shallow questions, but the largest portion of questions reflect some level of constructive problem-solving activity and show a positive relation with course grades. Questions that do not describe students' own reasoning or which do not display students' attempts to solve the problem seem to be related to students' prior experience. These constitute the smallest percentage of student questions. These

findings provide insight into the ways in which students benefit from engaging in online discussion forums.

2. RELATED WORK

The effectiveness of online discussions has been an area of active research for more than a decade. These results add to the long-standing body of work regarding how students learn CS: prior experience is influential [2, 18]; students’ expectations of their own performance are related to their achievement [14]; and lower-pressure environments may be more conducive to learning [19, 18].

The emerging body of empirical findings on discussion forums are wide ranging. There is evidence that students feel positively about discussion forums, believing those discussion forums benefit their learning. In a survey of 116 undergraduate and graduate students, students reported a greater sense of learning, convenience and flexibility, as well as the feeling of “less pressure” and being able to “openly express opinions” when they engaged in online discussions [19]. In a slightly more modest finding, Patel et al. [11] found that 55% of 62 undergraduate students surveyed felt that their learning in the class was helped by online discussions, while 45% reported feeling that the online discussions neither helped nor hindered their learning. Mihail et al. [8] also saw mixed results in a study of 43 students across two different on-campus computer science classes. Students in one computer science class perceived that posting in an online discussion helped their learning, while students in a different class did not.

There is evidence that forum participation and learning are related. Mihail et al. [8] found, in the same study mentioned above, that students who scored in the top third of the class posted more than students who scored in the bottom third of the class. Davis & Graff [5] studied 122 undergraduates, finding that among students who passed a class, there was no correlation between more online posting and final grade. However, students who failed the course interacted online less than other students. In other work that looked at a formally assessed discussion forum (in which participation was required for a course grade), most students posted only the minimum amount required to obtain course credit, and the number of new posts to the forum was positively associated with final grade [10].

Although MOOCs have important differences from in-person and blended courses, findings regarding discussion forums in MOOCs inform what we know about how students learn through online discussions. In MOOCs, students may have distinct styles of interaction with the MOOC as a whole and with the forums specifically [1]. In a computer science MOOC, a small fraction of users were responsible for a large portion of forum posts, and the activity on forums was positively associated with performance in the course. In analysis of a psychology MOOC, students who participated in the discussion forum had higher learning gains than those who did not. Furthermore, among those who did post, quantity of posts was positively associated with learning gain [17]. For the same MOOC, questions that demonstrated higher-order thinking were related to better outcomes [16].

Our work builds upon these empirical findings and relies on learning theories pertaining to the kind of active learning that discussion forums represent [4, 16]. Our sample of students is larger than the aforementioned in-person research studies and spans two different universities, while focusing

specifically on the CS2 course, which is a crucial stepping stone for students on their path to computer science degrees.

3. DATA COLLECTION

We collected Piazza logs from CS2 courses offered at North Carolina State University (University 1) and at the University of North Carolina (University 2). We analyzed Spring 2016 data for the two universities. At University 1 the CS2 enrollment was 214 and at University 2 the enrollment was 346.

Course Format. At both institutions, CS2 is taught in Java and programming projects are completed using the Eclipse IDE. All sections are taught in large lecture halls, and programming projects plus exams constitute the majority of the course grade. University 1 had 9 teaching assistants and University 2 had 17. With a total of 560 students and 26 TAs, the overall ratio of students to TAs is approximately 22:1.

Piazza Activity. All courses utilized Piazza, and 81% of enrolled students made at least one post on Piazza (post types include questions, notes, and answers). Nearly all students, over 99%, viewed Piazza posts. Across the two universities, we observed more than 250,000 views and 10,000 posts.

Consenting Students and Grades. For University 1, 66% ($n = 141$) of students consented to research, and at University 2, 90% consented ($n = 310$). Out of the students who consented, we analyzed those who had final grades, $n = 128$ for University 1 and $n = 267$ for University 2. The logs reflect 4,716 total questions, and we extracted only those questions from students who consented to this research, and for whom final grades were available. The resulting Piazza data set contains 1,679 questions from 395 students. On a scale of 100 points, the mean course grade at University 1 was 78.5, with a median of 87.8 ($\sigma = 25$). At University 2, the mean course grade was 83.1, with a median of 88.3 ($\sigma = 19$).

Students’ Prior Experience. We conducted a survey of consenting students that asked them to report their prior experience using Java. The options were as follows: high school classes; AP exam or class experience; professional experience; online course; self-taught; and experience from another college-level source. Students checked all options that applied. Figure 1 shows the reported experience by university. On the whole, 38% of students reported high school experience, 31% AP experience, 17% reported being self-taught, 22% had online course experience, and 7% reported professional experience with Java. 14% of students report other college-level Java experience from another source, which included a place-out exam for CS1, transfer credit for CS1, or an introductory course at another college or university.

Consistent with prior work [6, 9, 18], we found a relationship between students’ prior experience and their final course grade. We found that students who reported high school or professional experience tended to achieve higher grades in the course.

At University 1, the 47 students who reported high school experience had a significantly higher average grade of 82 ($\sigma = 23$), compared to an average of 72 ($\sigma = 29$) for the 46 students who did not report high school experience ($p = 0.0316$).¹ At University 2, there was no significant difference.

¹We consider p values less than 0.05 to be significant.

Table 1: Enrolled and Consenting Students

	Enrolled Students	Consenting Students	Consenting Students Assigned Final Grade	Questions from Consenting Students Assigned a Final Grade	Questions from Consenting Students with Final Grade who also Completed Experience Survey
Univ. 1	214	141	128	874	874
Univ. 2	346	310	267	968	805
Total	560	451	395	1842	1679

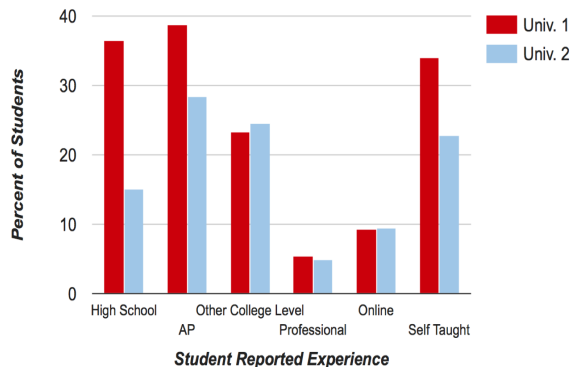


Figure 1: Students' experience percentages

The final grades of students at University 1 who reported professional experience ($n = 7, \mu = 91, \sigma = 13$) was significantly higher than those who did not ($n = 86, \mu = 76, \sigma = 27$), ($p = 0.014$). Once again, at University 2, no significant difference was found.

For all the other experience categories (AP, self, online, other college-level experience), as well as aggregating those with no experience versus some experience, there was no statistically significant difference in grades.

4. QUESTION TYPE ANALYSIS

To gain deeper insights into the ways in which students engaged on Piazza, we manually labeled each question in the Piazza data set (all questions asked by consenting students who had a final grade). These criteria resulted in 874 questions from University 1 and 805 from University 2, for a total of 1,679.²

In determining how to classify the Piazza questions, we drew on the question-asking literature as well as theories of learning. First, a rich history of literature has looked at classification schemes for questions and the frequency of the questions, finding that, for example, students tend to ask shallow questions overall. One category of question has been termed "knowledge deficit", through which students ask a question that stems from encountering an obstacle or contradiction [12]. The majority of questions in our Piazza data set represent this situation, as CS2 students often used Piazza to ask for help on programming projects.

In seminal work on how students engage in active learning, Chi [4] differentiates (from high to low levels of reasoning and connectedness) *Interactive*, *Constructive*, and *Active* engagement. Chi's framework is grounded in learning theory

²It is desirable to investigate answers and follow-on discussions to these questions. This paper leaves that task to future work (see Section 7).

on *constructivism*: rather than passively absorbing information, students *construct* the relationships they need through engaging actively. In this sense, all discussion forum posts are an example of active learning because, by definition, writing a post is not a passive activity. However, student posts vary widely in the level of reasoning that they reflect, and hence in whether they represent what Chi terms *Interactive* (high reasoning and connection to the contributions of others), *Constructive* (in the sense of constructivism: piecing ideas together), or at the lowest level, *Active* (posts that do not reflect personal reasoning) contributions. Note that *Active* only indicates that the student has engaged in the discussion forum. The Interactive-Constructive-Active framework has recently been applied in the context of MOOCs, where it was found that displaying higher-order thinking as reflected through interactive and constructive posts was associated with improved learning in the course [16].

We adapted the Interactive-Constructive-Active framework for our question classification scheme, with modifications based on the specific CS2 context. We separate questions about students' work from questions about course materials and policies. Questions about student work may be one of the following: *Constructive*, meaning that they display reasoning or a connecting sequence of students' attempts to solve the problem; or *Active*, if they provide minimal context for the problem, display a lack of personal reasoning towards solving the problem, and imply an open-ended request for help. (The *Interactive* tag occurred infrequently in our data and was combined into the *Constructive* tag.) In addition to these questions on student work, the data contain questions on the course itself. *Logistical* questions regard such matters as where to submit assignments or the schedule of exams. *Content-Clarification* questions ask for additional information on project assignments or auto-graders and are not in the context of a student's own work. The resulting question classification scheme is shown in Table 2.

To apply this classification scheme, we followed best practices to establish inter-rater reliability. One trained researcher labeled the entire data set, while a second trained researcher independently labeled a randomly selected 20% of questions, stratified by university. Cohen's Kappa, a measure of agreement that adjusts for agreement by chance, was $\kappa = 0.81$, indicating more than sufficient reliability to proceed with the following analyses [7].

5. RESULTS

We investigated the relationship between question asking and how students fared in the course as a whole. Using

Table 2: Question Classification Scheme

Type	Description	Example
<i>Active</i>	Request for help that does not display reasoning and is not framed in what the student knows or has already tried.	<i>My planets begin to orbit in the shape similar to the video, but then they stop moving. any way to fix this?</i>
<i>Constructive</i>	Questions that reflect students’ reasoning or attempts to construct a solution to the problem.	<i>How does one define “strongly typed?” based on that why is Java not a strongly typed language? I was under the impression that that meant each variable had a type set to it. Is java not strongly typed because you can typecast variables?</i>
<i>Logistical</i>	Questions on course policies, schedules, and assignment submission mechanisms. Not necessarily related to computer science content.	<i>On the syllabus, the office hours for TA/LA’s aren’t listed. Where can I find them?</i>
<i>Content-Clarification</i>	Request for additional information on project assignments, automated software tests, and design documents. Does not involve a question on the student’s own problem-solving work. (If a question does involve the student’s own work, that question would either be labeled <i>Active</i> or <i>Constructive</i> .)	<i>I’m a bit confused on what we are doing exactly. Are we asking the user to input all of the data and then after they’ve input all of the data, then we format it. Or are we making a program that formats data that has already been inputted?</i>

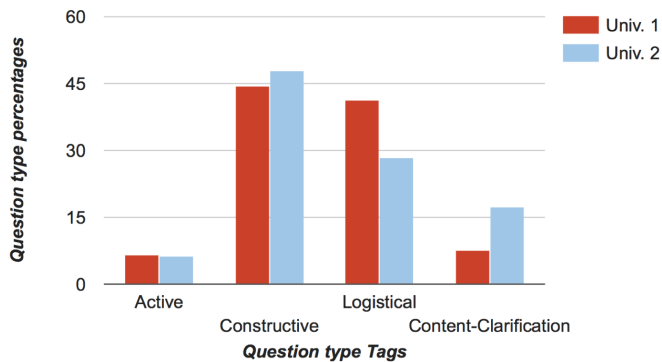


Figure 2: Question type percentage across schools

course grade information, we examined whether asking a question was related to course performance.

Students who asked questions on Piazza did indeed perform better in terms of final grade. At University 1, final grades of students who asked at least one question averaged 83 ($n = 105, \sigma = 20$), and students who did not ask any questions averaged a final grade of 58 ($n = 22, \sigma = 34$). This difference is statistically significant ($p = 0.001$). For students at University 2 who asked questions ($n=141$), the final average was 88 ($\sigma = 10$), compared to 83 ($\sigma = 17$) for those who did not ($n = 126$). This difference is also statistically significant ($p = 0.005$).

With the questions tagged, we proceeded to analyze their frequency distribution and to explore whether different types of questions were associated with students’ final grades in the course.

5.1 Distribution of Question Types

As shown in Figure 2, *Constructive* questions account for 46% of all questions. This finding is perhaps not surprising when we consider that solving programming projects is a highly constructive activity (in the literal sense of “constructing” a solution). Questions that ask for help with a

software bug, for example, are labeled *Constructive* if they provide some reasoning or connect together ways that the student (or others) have tried to solve the problem. In contrast to these questions, those that were labeled *Active* account for only 6% of all questions. Keep in mind that this label represents a question that reflects minimal reasoning, and many of them are shallow or provide very little context. *Logistical* questions (35% of all questions) and *Content-Clarification* questions (12% of all questions) account for the remaining posts in the question data set. Both of these question types ask about general class issues and do not involve a student’s own problem solving.

5.2 Constructive Questions

In order to more closely examine the benefits of asking *Constructive* type questions, we compared the set of students who asked at least one question of this type to students who did not. To investigate this, we first narrowed the data to students who posted at least one question of any kind throughout the semester. These are the “question askers”. Out of all question askers, the number of *Constructive* question askers varied across the two universities. At University 1, the percentage of *Constructive* questions were less prevalent, with 76% of question askers asking at least one (80/105). However, at University 2, only 67% of question askers (95/141) asked one or more *Constructive* questions.

Looking more closely at all *Constructive* question askers, 34% asked only one *Constructive* question during the academic term. Almost 58% of *Constructive* question-asking students posted 2-10 *Constructive* questions, and the number of students asking higher numbers of *Constructive* questions decreased from there (Figure 3). However, we did observe three outlying students who asked 20-75 *Constructive* questions, which are not shown in the figure.

Next, we examined the final grades of students who asked *Constructive* questions versus those who did not. In University 2, students who asked *Constructive* questions ($n = 95$) earned an average course grade of 90 ($\sigma = 9$), while students who asked questions but no *Constructive* questions ($n = 46$), achieved an average course grade of 85 ($\sigma = 11$). This dif-

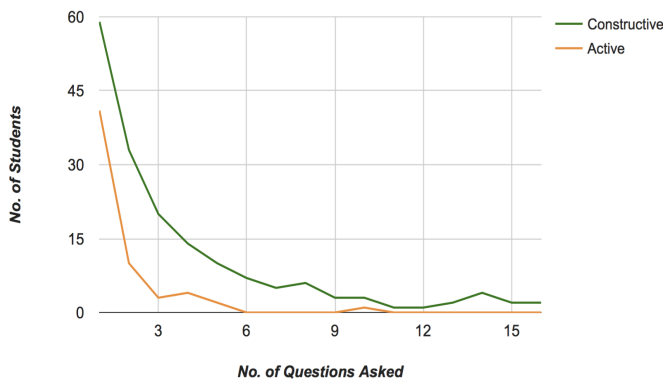


Figure 3: Question distribution per student

ference is significant ($p = 0.007$). There was no significant difference for University 1. We did not find any statistical significance in *Constructive* questions based on experience type.

5.3 Active Questions

Active questions, which do not reflect reasoning or an attempt to solve the problem on the part of the student, constituted 6% of questions overall. Through these questions, students ask for help without displaying their thought process, and without describing their previous attempts to solve the problem. Though these questions often do elicit helpful answers, they are unlikely to be as beneficial as *Constructive* questions.

Considering only students who posted at least one question throughout the semester, at University 1, the percentage of *Active* question askers was 27%. In contrast, at University 2, only 9% of students asked one or more *Active* questions.

Looking more closely at all *Active* question askers, the majority (67%) asked only one *Active* question during the academic term. Students rarely asked more than one *Active* question during the semester, and only one student asked as many as ten, as shown in Figure 3. In contrast to the finding with *Constructive* questions, there was no significant difference in the grades of students who asked *Active* questions versus those who asked questions but who did not ask *Active* questions.

There was a relationship between *Active* questions and student prior experience. For students who asked questions, *Active* questions accounted for a smaller proportion of questions posted by students who reported AP experience. Students who reported AP experience ($n = 86, \mu = 5\%, \sigma = 12$), asked proportionally fewer *Active* questions versus students who did not report AP experience ($n = 133, \mu = 9\%, \sigma = 22$), with $p = 0.0222$.

At University 1, for students who asked questions, the students who reported professional experience ($n = 4, \mu = 1\%, \sigma = 1$) asked fewer *Active* questions than those with no reported professional experience ($n = 71, \mu = 11\%, \sigma = 22$), $p = 0.00012$. The difference was not significant at University 2.

6. DISCUSSION

The results from this study of CS2 across two different large universities provide insight into how computer science

students leverage and benefit from online discussions as part of the learning process. First, from the data set of more than 1,500 Piazza questions, it is clear that CS2 students used the online discussion not only to pose general course-related questions (which we labeled *Logistical* and *Content-Clarification* questions), but very frequently to ask questions about their own problem-solving processes. Among those questions about students’ problem solving—which we labeled *Constructive* or *Active*—*Constructive* questions were far more common. This is an encouraging result because *Constructive* questions indicate that students are reasoning or connecting ideas in pursuit of their problem-solving goals. In a computer science class, these questions reflect the processes we *hope* are happening.

On the other hand, questions that do not reflect reasoning and that do not connect the student’s own attempt to solve the problem are the kinds of questions we generally hope to train our students not to ask. For example, one *Active* question is, “I created a helper method and wanted to sum up all the return values when I used the for loops in another method. How am I going to do so?” This question certainly reflects active learning, as the student has thought through a design and identified a goal. However, it lacks reasoning and does not describe what the student has already tried, what has been learned in class, or what others have posted on the forum. Therefore, it represents what we consider a minimal type of contribution in terms of active learning.

The importance of students making deeper reasoning or connections within online discussions is highlighted by the positive relationship between *Constructive* questions and course grade. This finding is consistent with very recent findings in MOOCs [16] (though for a different discipline, psychology), which showed improved outcomes for students who asked more *Constructive* and *Interactive* questions. Neither of these studies speaks to causality, so it is important for future work to investigate the ways in which supporting students’ *Constructive* or *Interactive* question asking may further improve their outcomes.

Although there were not enough *Interactive* questions to analyze this tag separately for inter-rater reliability (and hence we collapsed it with *Constructive*), it is informative to consider one example of a question labeled *Interactive* from the data set. A student asks, “In class we were basically advised to always use Linked Lists (...) for Queues, which I understand, since that’s much more efficient. However, we were also told that Array lists are generally used for Stacks, and I’m not clear on why. Wouldn’t it be more efficient to store a Stack as a linked list with only a front reference? ...” During question tagging, only five questions out of more than 1,500 were identified as containing this highest-level type of reasoning. This is important information for CS Ed practitioners to keep in mind: our discussion forums are clearly valuable for supporting students’ problem solving, but these online forums may not foster conceptual talk. The extent to which we could scaffold conceptual talk and deeper *Interactive* contributions is an important consideration.

7. CONCLUSION AND FUTURE WORK

This paper has presented an empirical analysis of discussion forum use in CS2, finding that students use the forum for a wide variety of course-related discussions. Many times, these discussions are not related to students’ own work, but are logistical in nature or to clarify the content of course

materials. However, when students do post questions about their own work, questions that reflect their reasoning are associated with better learning outcomes.

The findings, while not implying causality, suggest several considerations for CS faculty. First, faculty may consider encouraging students to actively engage in the class online discussion forum. In our data, students who posted at least one question did better in the course. Encouraging questions that reflect reasoning or that connect the student's attempts to solve the problem may be particularly helpful, as we found that these questions are associated with better outcomes. Second, care should be taken in responding to questions that do not reflect students' reasoning. Students with less depth of experience in computer science are more likely to ask these questions, and they may need scaffolding to learn how to better reflect their reasoning or what they have tried when asking a question. Finally, faculty may consider identifying students who can "seed" the forum with questions that bring multiple contributions together, as these higher-order questions may be particularly beneficial to foster in our online discussion forums.

There are many important directions for future work. First, students' characteristics, such as gender, ethnicity, personality, and attitudes likely interact in important ways with question types and benefits of discussion forum engagement. Future work should model this rich set of student characteristics. Second, analysis of formative grades, such as homework or projects, can provide a more sequential view into how discussion forum engagement impacts learning. Finally, we have examined questions, but each question post typically spawns a multi-turn answer and follow-on discussion that may yield great insights into how students are learning. By pursuing these lines of investigation, we can come to more clearly understand how to best support our computer science students through online resources.

Acknowledgements. The authors wish to thank the members of the LearnDialogue group for their helpful input. This work is supported in part by Google through a Capacity Research Award and by the National Science Foundation through grant CNS-1622438. Any opinions, findings, conclusions, or recommendations expressed in this report are those of the participants, and do not necessarily represent the official views, opinions, or policy of the National Science Foundation.

8. REFERENCES

- [1] A. Anderson, D. Huttenlocher, J. Kleinberg, and J. Leskovec. Engaging with massive online courses. *Proceedings of the 23rd International Conference on World Wide Web*, pages 687–698, 2014.
- [2] J. Bennedsen and M. E. Caspersen. An investigation of potential success factors for an introductory model-driven programming course. *Proceedings of the First International Workshop on Computing Education Research*, pages 155–163, 2005.
- [3] L. Breslow, D. E. Pritchard, J. DeBoer, G. S. Stump, A. D. Ho, and D. T. Seaton. Studying learning in the worldwide classroom: Research into edX's first MOOC. *Research & Practice in Assessment*, 8(March 2012):13–25, 2013.
- [4] M. T. H. Chi. Active-Constructive-Interactive: A Conceptual Framework for Differentiating Learning Activities. *Topics in Cognitive Science*, 1(1):73–105, 2009.
- [5] J. Davies and M. Graff. Performance in e-learning: online participation and student grades. *British Journal of Educational Technology*, 36(4):657–663, 2005.
- [6] E. Holden and E. Weeden. The impact of prior experience in an information technology programming course sequence. *Proceedings of the 4th conference on Information technology curriculum*, pages 41–46, 2003.
- [7] J. R. Landis and G. G. Koch. The Measurement of Observer Agreement for Categorical Data. *International Biometric Society*, 33(1):159–174, 2008.
- [8] R. P. Mihail, B. Rubin, and J. Goldsmith. Online Discussions: Improving Education in CS? *Proceedings of the 45th ACM Technical Symposium on Computer Science Education*, pages 409–414, 2014.
- [9] M. Morrison and T. S. Newman. A study of the impact of student background and preparedness on outcomes in CS I. *SIGCSE Bulletin*, 33(1):179–183, 2001.
- [10] S. Palmer, D. Holt, and S. Bray. Does the discussion help? The impact of a formally assessed online discussion on final student results. *British Journal of Educational Technology*, 39(5):847–858, 2008.
- [11] J. Patel and A. Aghayere. Students' Perspective on the Impact of a Web-based Discussion Forum on Student Learning. *Proceedings of the Frontiers in Education Conference*, pages 26–31, 2006.
- [12] N. K. Person, A. C. Graesser, J. P. Magliano, and R. J. Kreuz. Inferring what the student knows in one-to-one tutoring: The role of student questions and answers. *Learning and Individual Differences*, 6(2):205–229, 1994.
- [13] C. Rogerson and E. Scott. The Fear Factor : How It Affects Students Learning to Program in a Tertiary Environment. *Journal of Information Technology Education*, 9:147–171, 2010.
- [14] N. Rountree, J. Rountree, A. Robins, and R. Hannah. Interacting factors that predict success and failure in a CS1 course. *ACM SIGCSE Bulletin*, 36(4):101, 2004.
- [15] H. Van der Meij. Constraints on question asking in classrooms. *Journal of Educational Psychology*, 80(3):401–405, 1988.
- [16] X. Wang, M. Wen, and C. P. Rosé. Towards triggering higher-order thinking behaviors in MOOCs. *Proceedings of the 6th International Conference on Analytics and Knowledge*, pages 398–407, 2016.
- [17] X. Wang, D. Yang, M. Wen, K. Koedinger, and C. P. Rosé. Investigating how student's cognitive behavior in MOOC discussion forums affect learning gains. *Proceedings of the 8th International Conference on Educational Data Mining*, pages 226–233, 2015.
- [18] B. C. Wilson and S. Shrock. Contributing to success in an introductory computer science course: a study of twelve factors. *ACM SIGCSE Bulletin*, 33(1):184–188, 2001.
- [19] D. Wu and S. R. Hiltz. Predicting Learning from Asynchronous Online Discussions. *Journal of Asynchronous Learning Networks*, 8(2):139–152, 2004.