TE-2021-000353

Less is More: Students Skim Lengthy Online Textbooks

Chelsea L. Gordon, Roman Lysecky, *Senior Member, IEEE*, and Frank Vahid, *Member, IEEE*

***Abstract*—Computer science textbooks with lengthy text explanations of concepts are often considered thorough and rigorous, so lengthy textbooks (and class notes) are commonplace. Some, however, suggest text should be concise because people tend to skim lengthy text. This paper takes advantage of modern digital textbooks that measure reading time to examine reading rates for various text passage lengths. For a widely-used CS textbook written in a non-concise style, students read shorter passages (200 words or less) at about 200 words per minute, which is a typical rate. But for longer passages (600+ words), the rate increased to about 800 words per minute, suggesting skimming rather than reading. For another widely-used CS textbook, from the same publisher but written in a concise style with text passage sizes kept below 250 words, students spent more time (around 200 words per minute) reading the text passages, and their time spent was well-correlated with text length, suggesting students were carefully reading rather than skimming. Across three digital textbooks, the more interactive elements (e.g., integrated questions) that were included, the more time students spent reading the text between those activities. The conclusion is that to best educate students, authors of CS content should take the extra time needed to explain concepts more concisely – a case of "less is more" – and incorporate many active learning opportunities.**

***Index Terms*—Social and professional topics • Professional topics • Computing education • Computing education programs • Computer science education**

I. INTRODUCTION

COMPUTER science college courses usually assign lengthy textbooks and require students to spend a significant portion of each week reading textbook content. Research shows this approach leads to many students skimming the readings or skipping them entirely, hampering learning [1, 2].

One study [1] surveying students enrolled in various disciplines from two universities found that only 25% of students read assigned materials before coming to class. Only 40% of students completed the assigned readings in preparation for exams, and 19% reported not completing their reading assignments at all. Furthermore, 62% of students spent less than an hour *total* on assigned reading, while only 6% spent more than two hours, suggesting that even those who do read are skimming, rather than reading in full.

Reading textbooks increases students' learning of college course material and influences the grade a student will attain. One study [3] found in an analysis of 12 undergraduate introductory science courses, students who reported reading their textbooks *often* had better grades than students who reported only reading *sometimes*. The authors also found that students who reported reading rarely, surprisingly, did as well as those who read often. In other words, some students may not benefit from reading textbooks, but those who do benefit receive a greater benefit from reading more of the book. Another study [4] found that students in an engineering design course were more prepared and more sophisticated in their solutions for open-ended engineering design problems when they read the textbook. A relationship has also been found between the proportion of readings completed and quiz scores and final grades in introductory psychology courses [5].

Strategies that have been employed to increase student reading include giving pop quizzes, awarding extra credit for reading, providing extensive instructor feedback on assignments and exams, and using active learning questions [6]. Many of these strategies find success [7, 8, 9, 10, 11], but have not been comprehensively effective for encouraging all students to read content. In addition, students report *knowing* that it is important to read *and knowing* that reading will improve their final grades [12].

A strategy that has been less explored is providing less text for students to read. Importantly, less text does not necessarily mean less content. Rather, less text can mean text that is more concise and targeted (which can be *harder* to author), and less jargony and wandering. Using less text can reduce the sense of being overwhelmed that is associated with lengthy assigned readings.

For this study, student behavioral data was obtained from a major Computer Science textbook publisher that offers digital textbooks, some with more concise text than others. Those digital textbooks also contain animations and active learning questions that serve a related purpose to reading, namely to promote student learning.

The publisher's platform records student activity, which

Chelsea L. Gordon is with zyBooks, Campbell, CA 95008 USA (e-mail: chelsea.gordon@zybooks.com).
Roman Lysecky is with zyBooks, Campbell, CA 95008 USA and University of Arizona, Tucson, AZ 85721 USA (e-mail: roman.lysecky@zybooks.com).
Frank Vahid is with zyBooks, Campbell, CA 95008 USA and University of California, Riverside, CA 92521 (e-mail: Frank.Vahid@zybooks.com).

enables the tracking of time spent in the book overall, as well as time spent on particular kinds of activities (e.g., animations, learning questions, reading text). If all of the students were reading all of the text, one would expect to see two trends in the data. First, one would expect for textbook chapters that have more text to have substantially longer reading time from students. Second, one would find a correlation within the textbooks between the length of a text passage (in words) and the amount of time that students spend reading that passage.

Four hypotheses were proposed, namely that students would be more engaged with the reading, spending more time per word, for:

1. *Shorter text passages*, as students would see a short passage and decide to read it completely, while skimming longer passages.
2. *Shorter sections*, as students might start a section by scanning the entire section's length, then determining their reading pace based on available time.
3. *Shorter first text passages in a section*, as a longer first section might frustrate students who then decide to skim the rest of the section.
4. *Sections with more activities*, as by keeping students engage, students may read to be able to complete those activities.

II. Methods

This study looked at the student reading times for content in 12 courses using three different top-selling textbooks on the subjects of Python, Computer Organization and Systems (COS), and Discrete Mathematics (DM). Each book is used by over 30,000 students per year. Each college course using such a book has its activity recorded. Courses using those books were selected for this study if:

- the course had at least 50 students enrolled
- the instructor of the course had assigned readings for the course
- the course took place in Fall 2019 (pre-COVID).

Of the courses that met these criteria, four were selected at random for each subject area.

Table I provides passage and section length descriptive statistics for each book. For example, for COS, the average text passage was lengthy (mean 455 words, max 2900+). *Text passage length* is the number of words of text separating two activities. The length of an entire section (e.g., chapter 1, section 2) tended to be long (mean 2818 words, max 9000+). The textbook included interactive activities in the form of animations, formative learning questions, and auto-graded homework problems. A typical section had around 9 such activities, with a maximum of 15. The other books, Python and DM, were written natively for the web and thus have shorter passages and sections (focusing more on teaching through activities), with Python being the more concise of the two.

The average number of sentences in a paragraph and words in a sentence are used as an approximate measure for text complexity. COS has more sentences per paragraph and more words per sentence than the other books, indicating higher complexity. Microsoft Word's text analysis tool was used to calculate approximate grade level for the text and the Flesch Reading Ease score, which measures readability of text using the average sentence length and average syllables per word. The approximate grade level for COS was much higher than the others, and the COS Reading Ease score was much lower as well. Overall, this comparison shows that COS is a more difficult textbook to read. It is possible that because of this, students will be more likely to skim the text.

100 text passages each were analyzed from four Python courses, four DM courses, and four COS courses. A student's data was included in the analysis if that student had completed at least 40% of the measured passages.

TABLE I
Descriptive statistics for the 3 textbooks. μ= mean.

	Python	**Computer Organization & Systems**	**Discrete Mathematics**
Passage length	μ: 181 max: 380	μ:455 max:2900	μ: 207 max:901
Section length	μ: 475 max: 1400	μ: 3818 max: 9000	μ: 1025 max: 2253
Number of activities	μ: 7 max:14	μ: 9 max: 15	μ: 11 max: 18
Avg sentences per paragraph	2.2	3.5	2.8
Avg words per sentence	16.2	21.9	16.9
Flesch Reading Ease	55	39.3	65
Grade level	9.4	13	8.3
Passive sentences	31.1%	21.4%	32.5%
Students	1257	475	481

A. Metrics

1) Text passage length

Text passage length is measured as the total number of words in the text separating two activities.

2) Reading time

Reading time for a given text passage was calculated as the time of the last interaction with the activity that came before, subtracted from the time of the first interaction with the activity that came after. Thus, these times do not exactly correspond to the time that students spent reading the passage, but rather the reading time plus the time needed to read the next activity instructions and then interact (typically adding 20-30 sec). These reading times could only be collected if the student completed the activity before and after the passage. In order to control for cases of distraction, where students may leave and come back mid-passage, reading times that were more than two standard deviations from the average for that passage were excluded as outliers.

3) Entire section length

Entire section length is measured as the total number of words in all text passages in a given book section (e.g., section 5.1). There are multiple text passages per section.

4) Number of activities

Number of activities corresponds to the total number of activities (animations, formative learning questions, and auto-graded homework assignments) in a given section. There are multiple activities per section.

5) Number of graphs

Number of graphs was included to control for time spent looking at figures, which are not counted in the passage length.

6) Struggle

The proportion of auto-graded homework assignments that students struggled on was computed. An assignment was considered struggled on if a student took more than three attempts before solving the activity, and that number of attempts was at least two times that of the average student.

7) Reading Score

In order to measure reading behavior, a student's *reading score* was calculated. This score is the raw correlation between the time that a student spent with a text passage and the length of that text passage. The higher the correlation, the more it is assumed they were reading the text. An example of a strong correlation is shown in Figure 1.

8) Reading Rate

A student's reading rate is the average seconds per word (spw) at which the student read all text passages.

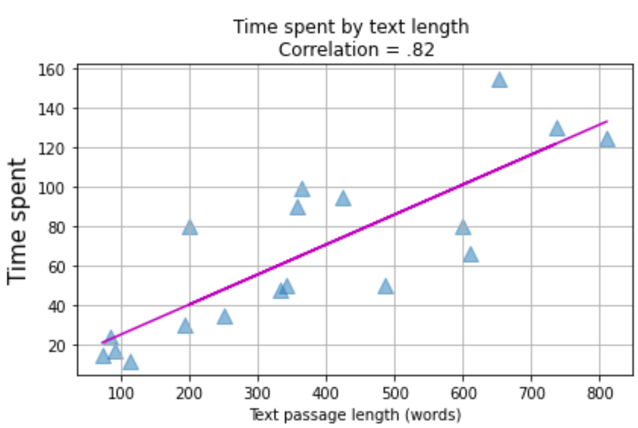

Fig. 1. Example reader score of 0.82, for sample times collected from the platform as one of this paper's authors diligently read text passages of varying lengths.

III. Analysis and Results

A linear mixed effects model was fitted to the data. The dependent variable was the log transformation of seconds spent per 100 words. This captures how engaged with the reading students were (more time means more engagement). The log transformation was used to normalize the distribution, as non-normality was observed in the raw data. Students were included as a random effect (intercepts only), to control for individual differences in reading speed. The number of graphs or figures included in a text passage was included as a control variable, to account for additional time that students spent viewing the graphs/figures. All predictor variables (fixed effects) were mean-centered and scaled, such that all values represent standard deviations from the mean estimate. The predictor variables included the:

1. length of the text passage
2. length of the total text in that section
3. length of the first text passage in that section
4. number of activities in that section
5. number of graphs in the text passage

To test for collinearity of the predictor variables, the variance inflation factor (VIF) was computed between all predictors. When the length of the first passage was included as a predictor, multiple variables exhibited collinearity factors of 4.0+, so the analysis was repeated without first passage, and all remaining predictor variables were found to have a VIF score of below 2.6, well below the recommended maximum VIF of 4.0. First passage length was thus not included as a predictor in the final model.

Likelihood ratio tests comparing models with the variables of interest to models without those variables were performed to obtain p-values. For Python, all predictors except for overall section length were significant at $p < .001$. For both Computer Organization and Systems and Discrete Mathematics, all predictors were significant. Full model results for each textbook are shown in Tables II-IV. *Estimates* refer to linear coefficients for a predictor while controlling for all other variables. Figures 2-5 below show the relationships between individual predictors and the time spent reading. Figure 6 shows the median seconds per word for each course in the analysis.

Figure 2 shows data for the first hypothesis, namely time spent per text passage by the text passage length. For a diligent reader, one would expect a strong positive correlation, namely that the longer the passage, the more time spent reading it. The "control group" line in the figure shows reading time by two people who were specifically asked to diligently read sections of various lengths from all three books. The average reading rate of control readers was 240 words per minute for COS, 182 words per minute for DM, and 188 words per minute for Python. The faster reading time for COS may be because the other two languages include more formulas and equations, which take additional time to comprehend. Python, with its short passages, exhibits a similar correlation from real students. However, the other two books show a positive relationship for passages under 300 words, but for longer passages the relationship disappears or even becomes negative, meaning students spent the same or less total time reading an 800 word passage as a 400 word passage.

Figure 3 shows the proportion of students skimming passages per passage length, irrespective of book. A skimmer is defined as someone reading 2.5x faster than the control group's average reading rate. For long passages of 400+ words, 40-60% of students skim.

Figure 4 shows data for the second hypothesis. As total section length increases, the reading rate worsens. The control group is the dashed line, which keeps a steady reading rate for any section length.

Support for the third hypothesis was not found, as first passage was collinear with other variables and was thus removed.

Figure 5 shows data for the fourth hypothesis. As the number of activities in a section increases, the reading rate improves for two books, but goes from bad to worse for the third book. This suggests students were skimming originally, and faced with numerous activities, skimmed even more, perhaps to focus on activities.

Lastly, an analysis was conducted to determine whether reader score predicted student struggle on auto-graded homework problems in the textbook. A generalized linear model was used with a binomial distribution of counts for assignments struggled on vs. assignments not struggled on. Struggle was regressed onto the reader score. Chi-square tests found that reader score significantly predicted struggle for both Discrete Mathematics and Computer Organization and Systems, (Chisq = 14.155, $p < .0001$; Chisq = 14.92, $p < .0001$, respectively), but not Python. Figure 7 shows these relationships.

A. Between-textbook reader score comparisons

The reader scores between the three textbooks were compared, with the expectation that reader scores for Python would be greater, because the text passages are shorter. The shapiro test showed no violation of normality, so a standard t-test was used to compare scores. Python and Discrete Mathematics both differed significantly from Computer Organization and Systems ($t = 17.6$, df 825, $p < .0001$; $t = 4.49$, df 985, $p < .0001$). Python and Discrete Mathematics were also found to be significantly different ($t = 14.7$, df 1168, $p < .0001$.

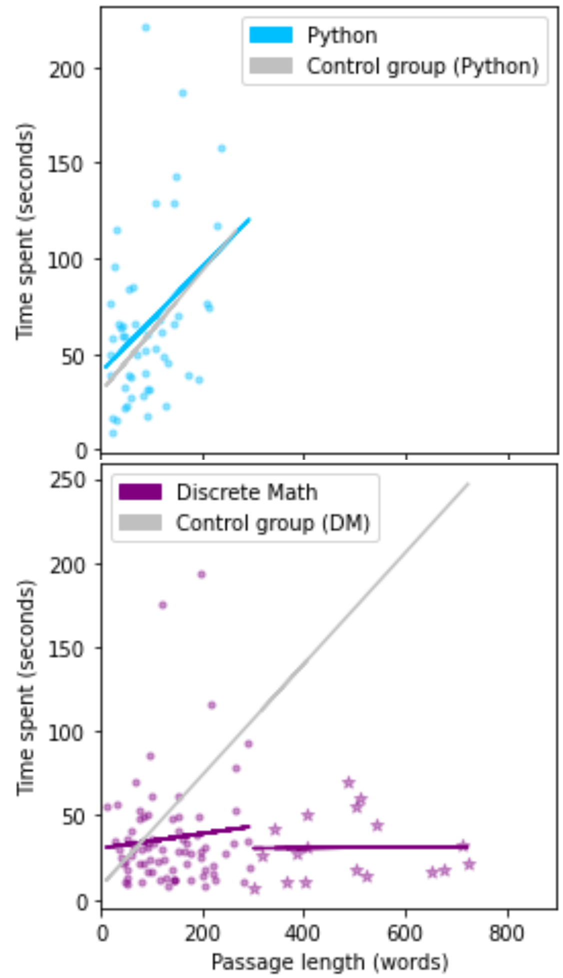

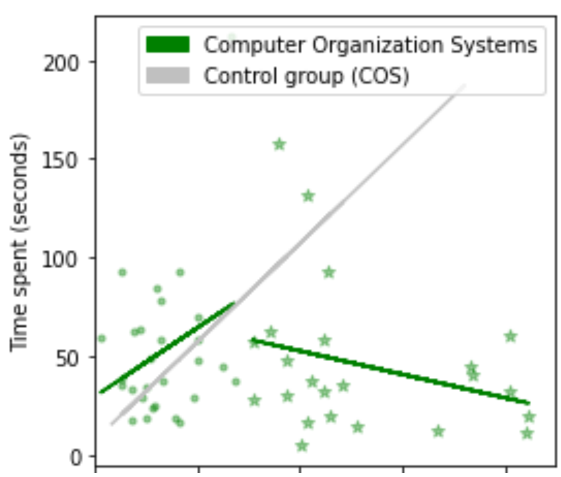

Fig. 2. Median time spent reading by passage length for students reading each of the three books, and for a control group. None of the sampled Python passages are above 250 words.

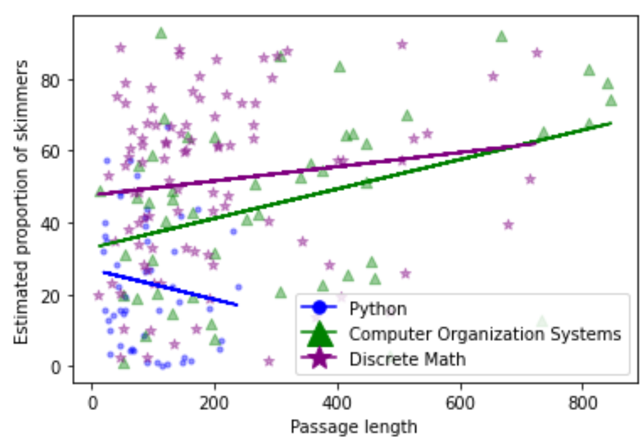

Fig. 3. Proportion of skimmers per passage length for each subject.

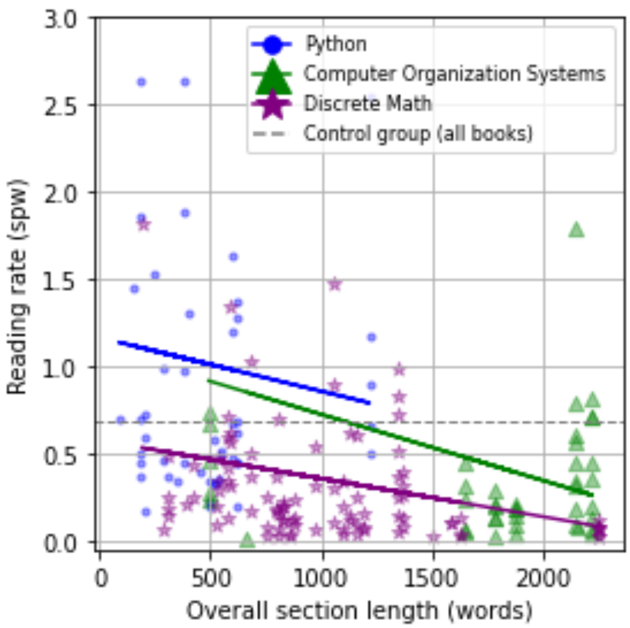

Fig. 4. Median reading rate for a text passage by total section length. As total section lengths increase, reading rate worsens. 4 points were above 3 spw and are omitted.

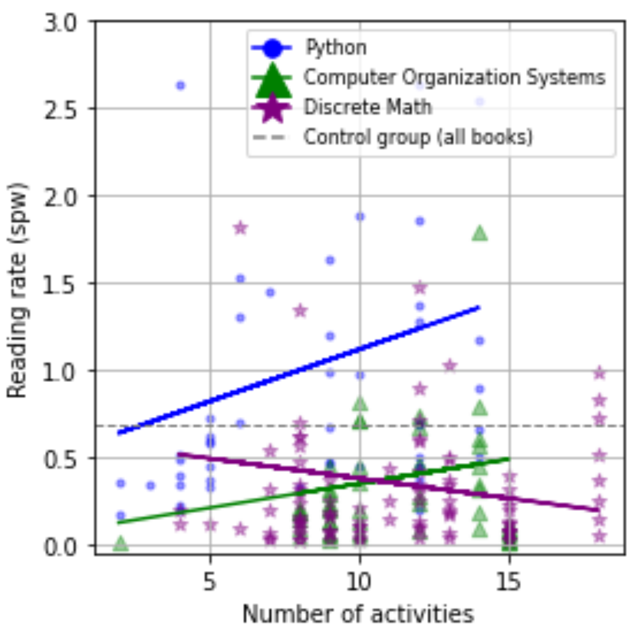

Fig. 5. Median reading rate for a text passage by number of activities in the section. As activities increase, time spent per word also increases. Four data points were above 3 seconds/word but are omitted on the y-axis.

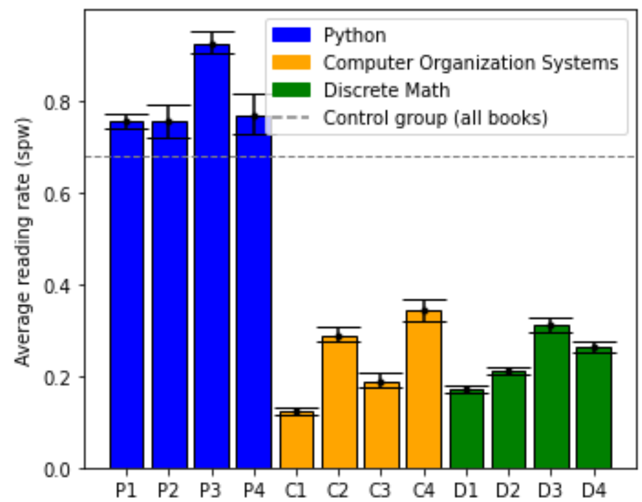

Fig. 6. Median reading rate for all 12 courses in the dataset, grouped by textbook. Python students average a bit higher than our control group average of 0.68, while those for Computer Organization and Systems and Discrete Mathematics are much lower, averaging around 0.2 seconds/word.

TABLE II

LINEAR MIXED EFFECTS MODEL PARAMETERS USING PYTHON COURSES. MARGINAL R^2 FOR ALL FIXED EFFECTS IS 0.12

Predictors	*Estimates*	*CI*	*p*	*df*
num_graphs	0.32	1.39 - 1.46	<0.001	45640
text_length	-0.70	-0.72 – -0.69	<0.001	45640
overall_section_text	0.01	-0.01 – 0.03	<0.433	45640
num_activities	0.12	0.10 – 0.13	<0.001	45640
Random Effects				
σ2	2.59			
τ00 user_id	0.29			
ICC	0.1			
N user_id	1312			
Observations	45647			
Marginal R^2 / Conditional R^2	0.120 / 0.208			

TABLE III

LINEAR MIXED EFFECTS MODEL PARAMETERS USING COMP ORG COURSES. MARGINAL R^2 FOR ALL FIXED EFFECTS IS 0.134

Predictors	*Estimates*	*CI*	*p*	*df*
num_graphs	0.11	0.08 - 0.13	<0.001	18729
text_length	-0.71	-0.73 – -0.69	<0.001	18729
overall_section_text	-.48	-0.50 – -0.46	<0.001	18729
num_activities	0.35	0.33 – 0.37	<0.001	18729
Random Effects				
σ2	1.08			
τ00 user_id	0.55			
ICC	0.34			
N user_id	495			
Observations	18729			
Marginal R^2 / Conditional R^2	0.344 / 0.565			

TABLE IV

LINEAR MIXED EFFECTS MODEL PARAMETERS USING DISCRETE MATH COURSES. MARGINAL R^2 FOR ALL FIXED EFFECTS IS 0.21

Predictors	*Estimates*	*CI*	*p*	*df*
(Intercept)	-1.47	-1.53 – -1.42	**<0.001**	32295.00
num_graphs	0.15	0.13 – 0.16	**<0.001**	32295.00
text_length	-0.66	-0.67 – -0.65	**<0.001**	32295.00
overall_section_text	-0.09	-0.10 – -0.07	**<0.001**	32295.00
num_activities	-0.08	-0.09 – -0.06	**<0.001**	32295.00
Random Effects				
σ^2	1.11			
$\tau_{00\ user\_id}$	0.49			
ICC	0.30			
N_{user_id}	581			
Observations	32302			
Marginal R^2 / Conditional R^2	0.217 / 0.455			

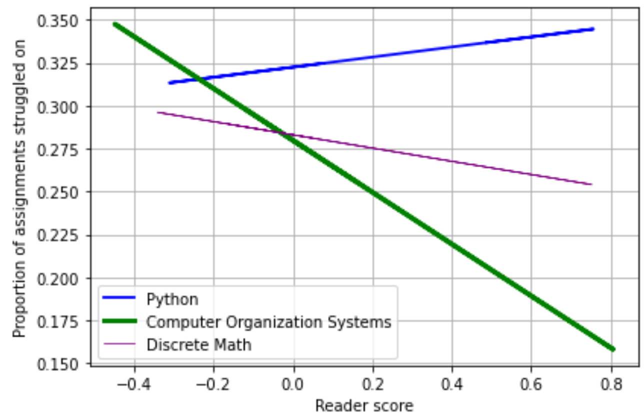

Fig. 7. Proportion of assignments struggled on vs. reader score. For DM and COS, where reading scores are lower on average, higher reading scores show lower struggle.

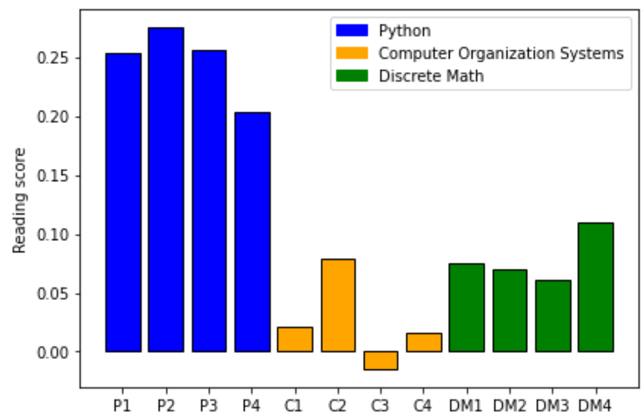

Fig. 8. Median reader score for all 12 courses. Python students' reader scores average around 0.25, while COS and DM were a much lower 0.04 and 0.10, respectively.

IV. DISCUSSION

Students read the Python textbook more diligently than the COS and DM textbooks, as seen by the comparison in reader scores as well as the reading rates in seconds per word. This difference might be because of the lengthier text in those books. However, another important difference is the type of course that uses the books. While the Python courses examined were mostly lower-division introductory courses, the COS and DM courses were a bit higher level courses, requiring students to have already completed some prerequisite courses. Therefore, another explanation for the different reading rates might be that first-year students are more likely to read textbooks (though

many would guess the opposite is true). Furthermore, as noted in Table 1, COS text had higher complexity and lower readability. The subject matter of the three textbooks also ranges from introductory programming to mathematics to computer architecture. As such, a direct comparison between these textbooks might not be as informative. Comparisons between more textbooks in future work will help distinguish between these hypotheses. This does not impact the finding that students in those latter classes were clearly skimming longer passages, which is still relevant to authors and teachers of such materials.

Whether the relationship between reading and struggle is causal is undetermined, as students who read more might also be students who struggle less by some unrelated factor, like better study habits or more commitment. Further controlled studies are needed to evaluate causality, but this study sheds light on the relationship between reading behavior and struggle, and potentially the importance of reading for student success. Additionally, the effect of reading behavior on struggle is not consistent across all three textbooks. COS showed the largest effect, while Python showed no reduced struggle with higher reading scores. Further studies should investigate what factors contribute to this relationship.

As first passage length was highly correlated with full section length, first passage length was not included in the mixed effects model. Therefore, no evidence shows that the length of the first section impacts reading rate as hypothesized. Further studies could investigate whether the first passage length does influence reading rate. However, these results suggest that keeping all passages consistently short encourages more diligent reading.

IV. Conclusions

These results provide evidence that using less text in introductory computer science textbooks can be an effective approach for increasing reading behavior in students. For three different textbook subjects, students were more willing to spend time reading text passages more carefully if those passages were shorter. In addition, students spent more time on individual text passages if a textbook section (e.g., chapter 3, section 1) had less text as a whole.

Reading long passages is mentally taxing, and college courses tend to have a lot of topics to read about. When passages are concise, students can quickly grasp new concepts before losing their attention, thereby leaving them to move on to the next topic without being stuck mid-understanding. Students also often peruse the length of an assigned section before beginning the reading. If a student sees that a book section is dauntingly long, the student may sometimes decide at that point that they will just skim this section, perhaps due to limited study time.

Students are willing to spend more time reading text if there are more interactive activities and examples within a section. When a textbook provides these hands-on experiences, students may be more engaged with the material, and their attention span is not limited to how long they can continuously consume text. Furthermore, such activities might motivate students to read the text to successfully complete those activities. If severely cutting away some of the text is not ideal, then chunking the text into smaller sections interspersed with engaging interactions can be effective.

Finally, students' reading behavior predicts how much they will struggle with the course content in Computer Organization and Systems and in Discrete Mathematics, here measured by number of attempts to solve auto-graded homework assignments, though the relationship may or may not be a causal one.

This research has several limitations and would benefit from additional research using complimentary methods. All the textbooks included in this study were written by the same authors in the same publishing company. As a result, they are written in a very particular style, which might not allow direct comparison to all textbooks, even within computer science. This research should be replicated on other textbooks from other publishers. An additional limitation is that the authors did not collect any qualitative data from students who were reading the books. Therefore, one cannot be sure that the longer reading times correspond to engagement; rather, they could possibly indicate confusion. Follow up studies might ask a subset of students to think aloud as they read through the text, giving insight into the student experience of short and long text passages.

While shorter text seems to encourage more reading in these textbooks, there is no evidence of a general rule across all disciplines and subject matter. The authors do not claim that long textbooks are generally problematic, and many systematic and comprehensive lengthy textbooks are highly praised. More systematic research can help discover which factors are important in the relationship between text length and reading behavior (e.g., introductory/advanced textbook, discipline and subject matter, course structure, price, level of interactivity, and more).

References

[1] K. Baier, C. Hendricks, W. Gorden, J. E. Hendricks and L. Cochran. “College students' textbook reading, or not,” *American reading forum annual yearbook,* 31, 385-402, 2011.

[2] M. A. Clump, H. Bauer, and C. Bradley. “The extent to which psychology students read textbooks: A multiple class analysis of reading across the psychology curriculum,” *Journal of Instructional Psychology*, 31, 3, 2004.

[3] M. French, T. Franco, M. Neumann, L. P. Kushnir, J. Harlow, D. Harrison, and R. Serbanescu. “Textbook use in the sciences and its relation to course performance,” *College Teaching*, 63, 4, 171-177, 2015.

[4] C. Atman and K. M. Bursic. “Teaching engineering design: Can reading a textbook make a difference?” In *Research in Engineering Design*, 8, 4, 240-250, 1996.

[5] E R. Landrum, RA.R. Gurung, and N Spann. “Assessments of textbook usage and the relationship to student course performance,” *College Teaching*, 60, 1, 17-24, 2012.

[6] S.A. Lei, K. A. Bartlett, and S. E. Gorney. “Resistance to reading compliance among college students: instructor's perspectives,” *College Student Journal*, 44, 2, 2010.

[7] B. Fleck, A. S. Richmond, H. M. Rauer, L. Beckman, and A. Lee. “Active reading questions as a strategy to support college students' textbook reading,” *Scholarship of Teaching and Learning in Psychology*, *3*(3), 220, 2017.

[8] D. McKinney, A.D. Edgcomb, R. Lysecky, and F. Vahid. “Improving pass rates by switching from a passive to an active learning textbook in CS0,” In the *Proceedings of the ASEE Annual Virtual Meeting*, 2020.

[9] A. Edgcomb, F. Vahid, R. Lysecky, and S. Lysecky. “Getting students to earnestly do reading, studying, and homework in an introductory programming class,” In *ACM SIGCSE Technical Symposium on Computer Science Education, Seattle WA,* 2017.

[10] A. Edgcomb, F. Vahid, R. Lysecky, A. Knoesen, R. Amirtharajah, and M. L. Dorf. “Student performance improvement using interactive textbooks: A three-university cross-semester analysis.” In *Proceedings of ASEE Annual Conference, Seattle WA,* 2015.

[11] A. Edgcomb and F. Vahid. "Effectiveness of online textbooks vs. interactive web-native content." In *Proceedings of ASEE Annual Conference, Seattle WA,* 2014.
[12] T. Berry, L. Cook, N. Hill, and K. Stevens. "An exploratory analysis of textbook usage and study habits: Misperceptions and barriers to success," *College Teaching*, 59, 1, 31-39, 2010.