Assessing the Need of Augmenting Video Lectures with Supporting Information

Gaurav Kumar Singh, Abhay Doke, Varun Kumar,
Savita Bhat, and Niranjan Pedanekar

Systems Research Lab, Tata Research Development and Design Centre,
Tata Consultancy Services, Pune, India
{gauravk.singh2,abhay.doke,kumar.varun1,
savita.bhat,n.pedanekar}@tcs.com

Abstract. Massively Online Open Courses (MOOCs) consist of online video lectures delivered by experts. Learner drop-out is a major concern for MOOCs. Early drop-outs are often associated with cognitive overload partially caused by unfamiliarity of concepts being taught. In such cases, the course can be augmented with supporting information such as definition and explanation for concepts. In this paper, we propose a metric quantifying the need for augmentation of individual concepts as a course progresses. We examine the metric using a MOOC course. We also present a preliminary experiment with 36 undergraduate students on using such augmentation.

Keywords: MOOCs, Education, Augmentation, Metric, e-learning.

1 Introduction

Over the last decade, Massively Online Open Courses (MOOCs) have gained popularity amongst learners. MOOCs typically consist of a number of video lectures of varying length grouped according to topics and arranged in an appropriate sequence. In a MOOC, the learners see a teacher, typically an expert in the field, delivering a lecture in front of a camera. The teacher can also present slides on the topic or write on a blackboard or pose questions which need to be answered by the learner. The learners can rewind, fast-forward or pause the videos according to their need.

However, MOOCs have failed to deliver on their disruptive promise mainly due to high attrition rates [1]. Studies indicate that over 90% of learners registering for MOOCs drop out without completing the course [2]. A study of the edX Circuits and Electronics course reported an attrition of 95% with almost 50% taking place in the early stages of the course [3].

One can argue that if the drop-outs could be prevented, MOOCs could have more business and social impact. In a review of the factors contributing to early drop-outs of learners in e-learning, Tyler-Smith [4] argued that the early drop-outs typically occurred due to a cognitive overload experienced by the learners. First-time learners deal with multiple tasks contributing to the cognitive load: being able to adopt tech-

P. Zaphiris and A. Ioannou (Eds.): LCT 2014, Part I, LNCS 8523, pp. 238–249, 2014.

nology, using the learning interface, taking on new concepts and interacting with other learners. In another work on MOOCs, Adamapoulos [5] reported that the conceptual difficulty experienced by a learner has a negative effect on the course completion rate.

Cognitive load may also be caused by a phenomenon known as the 'Curse of Knowledge' (COK) [6]. Due to COK, experts often tend to overlook the perspective of the novice and end up using unfamiliar or unrelated terms while teaching the main concept. For example, if a teacher expert in Unix/Linux casually says that 'information retrieval is a little like the *grep* command' in an information retrieval course, some learners may not understand what she wants to say. In an online setting, since the learners cannot ask questions, they have to pause the lecture and look up the unfamiliar term on the Internet. This involves hunting for information and trying to make sense of it in the context of what is being taught. This may further add to the learner's cognitive load.

Agrawal et al used various spatial and semantic characteristics of the learning material to propose a method for assessing comprehension burden in a textbook [7]. In another work, Agrawal et al [8] proposed augmenting textbooks with supporting information to ease the comprehension burden. One can argue that just like textbooks, if MOOC learners are provided supporting information for simplification of unfamiliar or important concepts at early stages of learning, they can counter the cognitive load to some extent. In fact, some studies [9] advocate the use of additional information from online resources while learning from MOOCs.

In a typical MOOC, thousands of concepts appear throughout the course. But which of these concepts should one augment? There seems to be no prior work on prescribing when augmentation should be provided during a MOOC to ease the burden on the learner. In this paper, we present, to our knowledge, the first attempt **to measure the augmentation need (AN) for concepts in a MOOC**. We specifically choose the familiarity of a concept to the learner as a basis for evaluating this need. We propose a metric for quantifying the augmentation need to determine whether a concept being mentioned needs to be augmented with supporting learning material at a given time during the course. In the proposed augmentation need metric, we incorporate the effects of the learner's familiarity with the concept, the importance of the concept in a particular lecture and the progress through the course. We report an analysis of a MOOC for Computer Architecture using this metric.

But does augmentation with additional information work towards reducing the cognitive load? We present a preliminary controlled experiment with undergraduate students in an attempt to answer the following questions:

1. Do learners of MOOCs need augmentation?
2. From a Human Computer Interface (HCI) perspective, how do first time learners receive such augmentation?
3. What effect does such augmentation have on the learner's understanding?

2 Augmenting MOOCs

One can imagine several dimensions of augmentation when it comes to MOOCs. We argue that a MOOC could be augmented with additional information such as: definitions to make learners familiar with a concept, explanations to make them understand concepts, complementary information (*e.g.* practical examples when theory is being explained), engagement mechanisms (*e.g.* rewards and challenges) when the lecture sounds monotonous, and assessment questions when a concept has been explained. In this paper, we focused only on the familiarity of a concept when assessing the augmentation need. For example, when a concept such as *grep* appears in an information retrieval course, it may not be very familiar to most learners. Perhaps a definition and an example of *grep* would help in making the learner understand the concept.

We make three main observations related to the need for augmentation of a concept during a lecture.

2.1 The Effect of Familiarity of a Concept

There are certain concepts in the domain of computer science that are more familiar to learners. A concept such as *Microsoft Windows* being mentioned by an instructor hardly needs any supporting information. A concept such as *bit* perhaps needs some introduction to the uninitiated. Advanced concepts such as *Re-order Buffer* need much more supporting information.

2.2 The Effect of Progress in a Course

As a course progresses, the learner is likely to know more about a concept. The need of the learner to make use of explanatory material may diminish over time. As the concept is mentioned more number of times, the augmentation need goes down.

2.3 The Effect of Importance of a Concept

A concept such as *sparse matrix* is important when describing the data structures used for information retrieval. But the concept may not even appear in other lectures of the course. So the need for supporting information for such concepts may be localized. A concept such as *corpus* may appear in many places in a natural language processing course and needs to be understood clearly.

3 Augmentation Need (AN) Metric

MOOC video lectures typically have subtitles. These are provided by the instructor or are available through crowdsourcing. We used these subtitles as the base for finding out what concepts are being mentioned by the instructor. We extracted the concepts from the course subtitles that have a corresponding Wikipedia page using the Wikipedia Miner [10] which allowed an accuracy of 75% in identifying Wikipedia links for given text [11]. The extracted concepts formed a set of augmentation candidate concepts.

Based on the three observations in previous section, we proposed a metric to quantify the need for augmentation of any given concept at any point during a course. We proposed that for a given concept c_i appearing in the course in lecture l, AN is a weighted average of the three effects mentioned above as

$$AN_{i,l} = \begin{cases} w_1 F_i + w_2 P_{i,l} + w_3 I_{i,l}, & c_i \text{ appears in } l \\ 0, & c_i \text{ does not appear in } l \end{cases}$$

where w_1, w_2 and w_3 are weights of the familiarity, progress and importance effects.

3.1 Familiarity Effect

For representing the familiarity of a concept, we proposed a metric based on the number of pages that refer to the concept's page in Wikipedia. For this, we formed a graph of all Wikipedia pages. We calculated the Global Familiarity (GF) for a concept c_i as

$$GF(i) = \log \frac{w_i}{W}$$

where w_i is the number of pages that link to the Wikipedia page for c_i and W is the total number of links in Wikipedia. A larger value of GF indicates a more familiar concept. For example, GF(*Microsoft Windows*) is -7.03, while GF(*Re-order Buffer*) is -14.11.

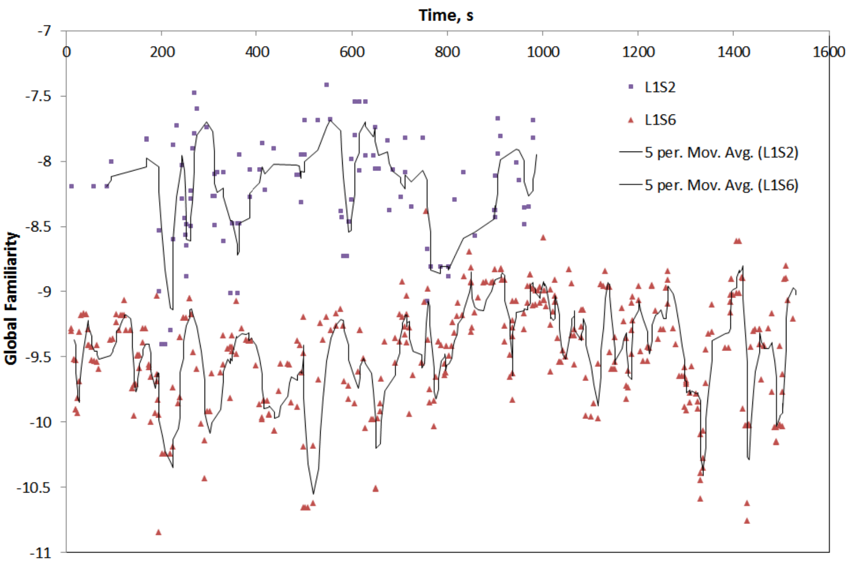

Fig. 1. Variation of Global Familiarity values and their 5-point moving averages for L1S2 (introductory) and L1S6 (advanced) lectures in the CA course

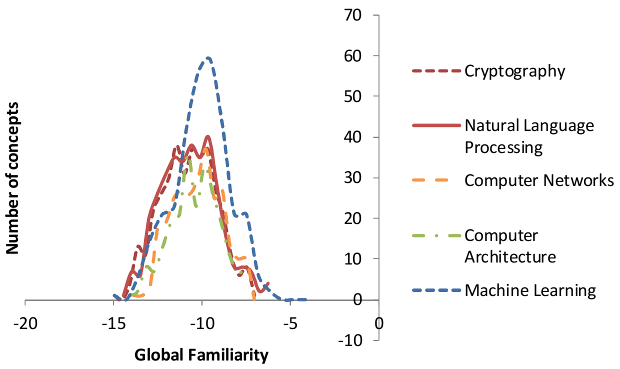

Fig. 2. Distribution of global familiarity over 5 MOOCs from Coursera

Figure 1 shows higher GF values for concepts occurring in an introductory lecture as compared to those in a more advanced lecture. We found GF scores for concept terms in 5 MOOCs from Coursera [12-16] and observed that the GF values are distributed in a similar way for all the courses as shown in Figure 2. From this we concluded that GF values are a consistent measure of familiarity of concepts. We normalized the GF for the concept c_i as

$$\overline{GF}(i) = \frac{GF(i) - GF_{min}}{GF_{max} - GF_{min}}$$

where GF_{max} and GF_{min} are GF values of the most familiar and the least familiar concepts in the course, respectively. We observed that very familiar terms do not need augmentation and they need to be given less importance in the metric. So we further shaped the normalized values as an inverted sigmoid curve as shown in Figure 3 for obtaining the global familiarity effect F_i for a concept c_i in a course using

$$F_i = \left(\frac{e^{-v}}{1 + e^{-v}}\right) \text{ where } v = 10(\overline{GF}(i) - 0.5)$$

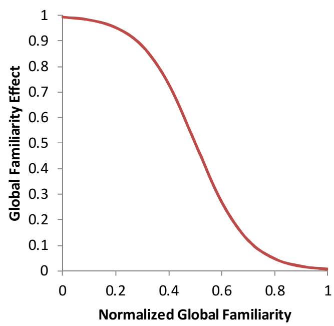

Fig. 3. Variation of familiarity effect F_i with the normalized global familiarity scores

3.2 Progress Effect

For representing progress through the course, we considered the fraction of a concept that remained to be talked about at a given point in the course. We proposed a metric to denote the progress effect P_i for a concept c_i as

$$P_i = 1 - \frac{n_{i,l-1}}{N_i}$$

where $n_{i,l-1}$ is the number of occurrences of c_i till the beginning of the lecture l and N_i is the number of occurrences of c_i in the whole course. The metric has an automatic value of 1 for the first lecture as no concept has been talked about yet. Figure 4 shows the variation of the progress effect with the frequency of a concept recorded over 8 lectures.

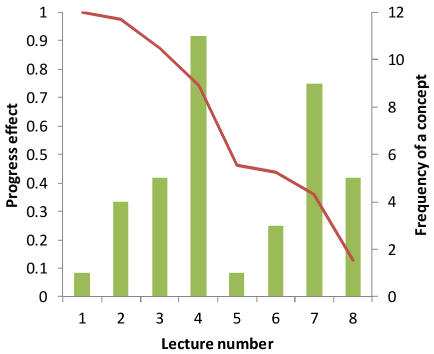

Fig. 4. Variation of the progress effect over lectures (shown as the line) for a given distribution of concept occurrences in those lectures (shown as bars)

3.3 Importance Effect

For incorporating the importance of concepts in a given lecture, we used the Term Frequency-Inverse Document Frequency measure (TF-IDF) as it is the standard measure of importance used in information retrieval [17]. We first calculated the importance $imp_{i,l}$ of a concept c_i in a given lecture l as

$$imp_{i,l} = f_{i,l} * log\left(\frac{L}{L_i}\right)$$

where $f_{i,l}$ is the frequency of occurrences of the concept c_i in a lecture l, L_i is the number of lectures in which the concept c_i appears and L is the total number of lectures.

We further normalized the importance values over each lecture to get the importance effect $I_{i,l}$ for a concept c_i in a lecture l as

$$I_{i,l} = \frac{imp_{i,l} - imp_{l,min}}{imp_{l,max} - imp_{l,min}}$$

where $imp_{l,min}$ and $imp_{l,max}$ are the minimum and maximum values of importance in lecture l, respectively.

4 Analyzing MOOCs Using the AN Metric

We analyzed a part of the Computer Architecture MOOC course offered by Prof. David Wentzlaff at Coursera [12] using the AN metric. This part has 27 lectures prescribed for a period of 6 weeks. Using the subtitles of the videos, we found 235 augmentation candidate concepts using Wikipedia Miner. We also recorded the time at which these concepts appeared in each lecture. We calculated the AN metric for each of these concepts in each lecture. We used equal weights in the weighted average in the AN metric equation (w_1 = 0.33, w_2 = 0.33 and w_3 = 0.33). We recommended candidate concept terms with AN values higher than a threshold of 0.33 amounting to at least one effect fully contributing to the AN value. This was done in order to suppress the number of augmentation candidate terms.

4.1 Observations

Figure 5 shows a partial visualization of the AN values for all 27 lectures in the Computer Architecture (CA) course. This visualization shows the top 10 concepts used in the course based on the average of their AN values over all lectures. Each column represents a lecture and each row represents a concept. The darker the cell colour, the greater is the augmentation need. The line graph at the end shows the variation of the AN value over the lectures. We observed that not all concepts are prescribed augmentation for all lectures. Some concepts such as *register file* are important throughout the course. These may need augmentation relevant to the context as the course proceeds. Concepts such as *Computer Architecture* appear early on, but are not prescribed for augmentation during later lectures. Concepts such as *Microsoft Windows*, *Linux* and *Java* (not seen in the figure) are not at all prescribed for augmentation.

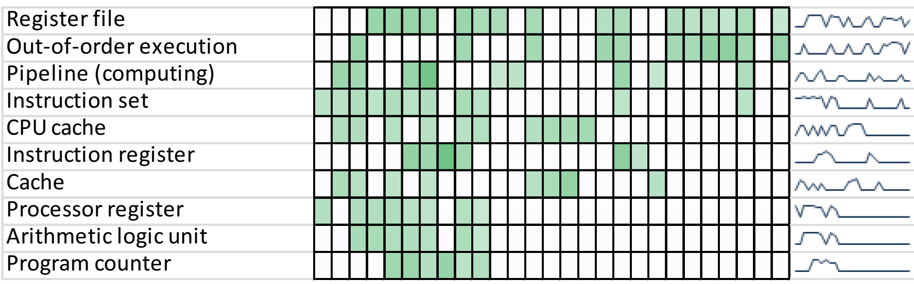

Fig. 5. Concepts with highest average AN values over the whole course

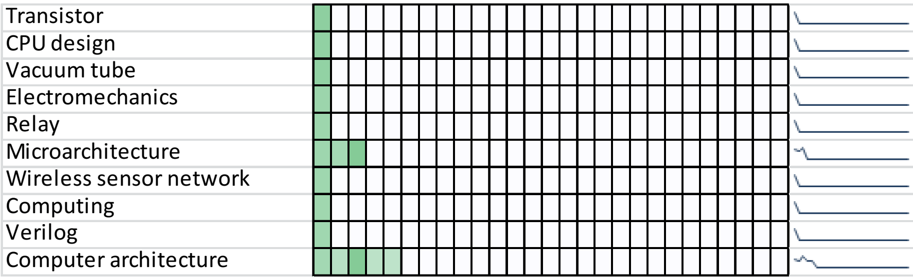

(a) Top terms for Week 1 – Introductory lecture

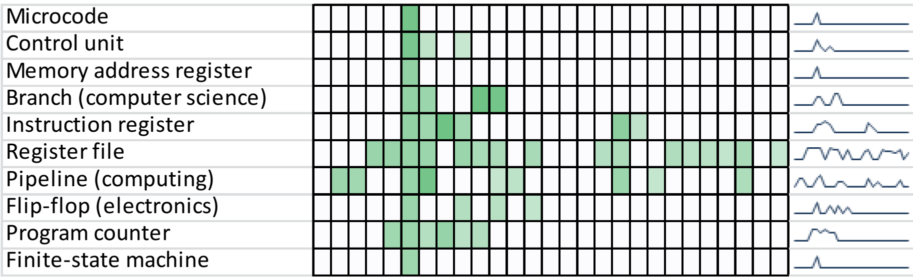

(b) Top terms for Week 2 – *Microcoded Architecture* lecture

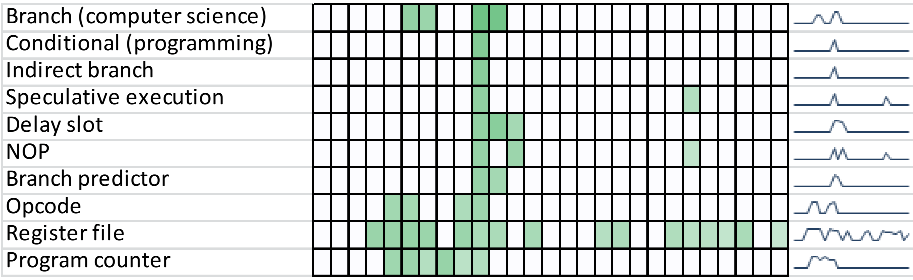

(c) Top terms for Week 3 – *Jumps* lecture

Fig. 6. Variation in the AN of concepts according to specific lectures

Also, for a given lecture the augmentation plan can be different than other lectures. For example, as seen in Figure 6(a), concepts such as *Moore's law* are prescribed for augmentation in the first introductory lecture. Most of these concepts may not be revisited later, but can cause discomfort early on if not familiarized. Figure 6(b) shows the top concepts in the first lecture of Week 2 about Microcoded Architecture of the CPU. Concepts such as *microcode* and *control unit* are important to this lecture and are not very familiar concepts. Therefore, they need to be augmented. Similarly, as seen in Figure 6(c), concepts such as *branching* and *conditional branching* are central to the first lecture in Week 3 about Jumps, and need to be augmented.

We also analyzed the relative contribution of the three effects on the AN value. As seen in Figure 7(a), a concept such as *pipelining* is not very familiar, but is very important in certain lectures. Though the progress effect diminishes later, the importance of the concept in the later lectures as well as the relative unfamiliarity of the concept are high. This causes the augmentation need to increase above the threshold for all lectures where *pipelining* appears. As seen in Figure (b), *bit* is not very important in lectures and is relatively familiar. So its augmentation need is below threshold for most lectures.

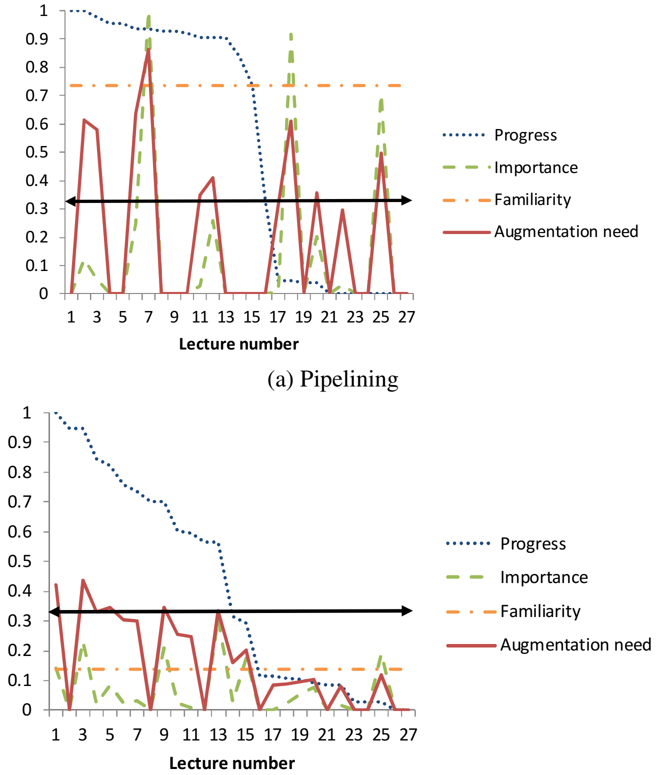

(a) Pipelining

(b) Bit

Fig. 7. Contribution of the three effects to the AN metric over the entire MOOC

5 A Preliminary Experiment with Augmentation

In earlier sections, we described the need for augmentation and proposed a metric for augmentation need specifically to avoid discomfort owing to unfamiliarity with concepts. In order to test whether such augmentation works in practice, we carried out a preliminary experiment with 36 undergraduate students studying Information Technology. We designed an augmentation web interface for one of the video lectures in the Natural Language Processing course by Jurafsky and Manning [13], which described the concept of term-document matrices used in information retrieval. The augmentation interface consisted of the video lecture accompanied by clickable buttons with concept names as shown in Figure 8.

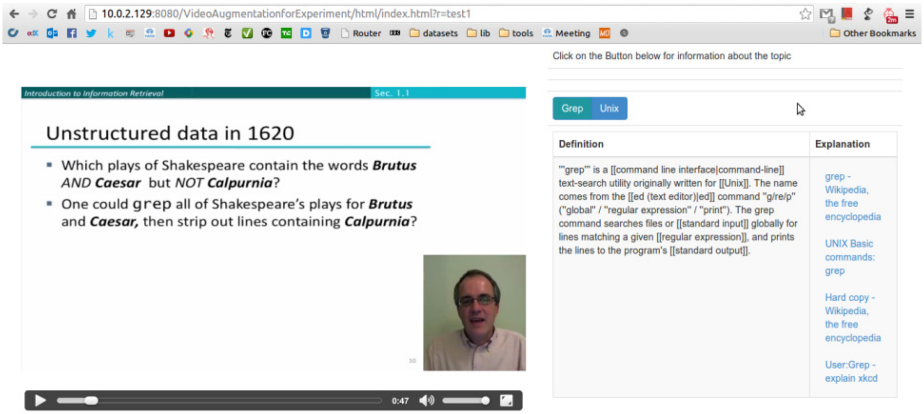

Fig. 8. Augmentation interface for a MOOC lecture used for the experiment

The buttons appeared in the side frame as a concept was mentioned, stayed for 10 seconds and faded away. Clicking the button paused the video and brought up preliminary information such as definition and links for explanation. The learner could use it if required and get back to the video by un-pausing it. We gave the augmented video to a test group of 19 students and the non-augmented video to a control group of 17 students. At the end of the 10-minute lecture, the learners appeared for a quiz consisting of 6 questions. They also answered survey questions about the need for, the usefulness and interestingness of, and the distraction due to augmentation. They also stated which concepts they found difficult. We also recorded the clicks of the learners as they used the augmentation information.

We found that there was no significant difference in the scores of the control and the test group. Also there was no significant difference in the perceived interestingness and degree of difficulty of the video between the control and the test group. 31% of the test group respondents said that they found the augmentation distracting, while 89% found it to be useful. 76% of the control group respondents said that they wanted additional information during the video. We also found that all the concepts which learners had stated to be difficult (*corpus*, *Caeser*, *grep*, *matrix, sparse matrix*, *Unix*, *data structure*, *string search*, *Boolean algebra*) were suggested as augmentation candidates using the AN metric. We observed that though learners in the test group had clicked on the buttons for definition, they did not extensively utilize the explanation links provided.

6 Discussion

Analysis of the Computer Architecture MOOC and the experimental results indicate that the AN metric works fairly well in identifying concepts needing supporting information. It captures the effects of familiarity, progress and importance on the augmentation need and reduces the set of candidate concepts to a manageable set for practical use. Due to the interplay of importance and progress effects, locally important terms which have to be understood well are given due importance. Globally importance terms also are recognized for augmentation, but the augmentation need diminishes over time. Very familiar terms do not receive augmentation, while the ones which are unfamiliar do get recommended for augmentation. One can vary the weights of the metric to suit the course needs, *e.g.* reduce the familiarity weight in an advanced course that contains difficult terms, so that even the terms which are less difficult get augmentation.

Among the main limitations of this approach, is the use of Wikipedia to estimate global familiarity. Concepts for which the global similarity value is not appropriate, the AN metric also does not make sense. We recommend using a more robust method for assigning global familiarity. Similarly, the use of Wikipedia Miner for recognizing concepts is another limitation as inaccurate recognition leads to redundant or incorrect concept candidates.

The results of the preliminary experiment suggest that learners needed augmentation, but did not specifically use the augmentation interface extensively to reduce their cognitive load. So despite the augmentation need being recognized, designing an interface that is intuitive to use and does not distract learners may be the key to reducing cognitive load.

7 Conclusion

In this paper, we presented, to our knowledge, the first attempt at assessing whether concepts being taught in a MOOC need supporting information. We proposed an augmentation need metric based on the effects of the familiarity of concepts, the progress in a course and the relative importance of a concept in a lecture. We also presented an analysis of applying the AN metric to a Computer Architecture MOOC. We reported experimental results on using augmentation interface with MOOCs. We believe that learners need augmentation during MOOCs and the AN metric provides MOOC instructors or designers a plan for using supporting information about concepts. We believe that we have demonstrated the utility of the AN metric, but a large scale validation is needed to assess it better. We plan to apply the AN metric to a large number of MOOCs in order to visualize, compare and contrast course characteristics.

References

1. Yang, D., Sinha, T., Adamson, D., Rose, C.P.: “Turn on, Tune in, Drop out”: Anticipating student dropouts in Massive Open Online Courses. In: NIPS Workshop on Data Driven Education (2013)
2. Jordan, K.: MOOC Completion Rates: The Data, `http://www.katyjordan.com/MOOCproject.html` (retrieved February 21, 2014)
3. Breslow, L.B., Pritchard, D.E., DeBoer, J., Stump, G.S., Ho, A.D., Seaton, D.T.: Studying learning in the worldwide classroom: Research into edX’s first MOOC. Research & Practice in Assessment 8, 13–25 (2013)

4. Tyler-Smith, K.: Early attrition among first time eLearners: A review of factors that contribute to drop-out, withdrawal and non-completion rates of adult learners undertaking eLearning programmes. Journal of Online learning and Teaching 2(2), 73–85 (2006)
5. Adamopoulos, P.: What Makes a Great MOOC? An Interdisciplinary Analysis of Student Retention in Online Courses. In: Proceedings of the 34th International Conference on Information Systems, ICIS, vol. 2013 (2013)
6. Wieman, C.E.: APS News–The back page. The "curse of knowledge" or why intuition about teaching often fails. American Physical Society News 16(10) (2007)
7. Agrawal, R., Chakraborty, S., Gollapudi, S., Kannan, A., Kenthapadi, K.: Empowering authors to diagnose comprehension burden in textbooks. In: Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 967–975. ACM (2012)
8. Agrawal, R., Gollapudi, S., Kenthapadi, K., Srivastava, N., Velu, R.: Enriching textbooks through data mining. In: Proceedings of the First ACM Symposium on Computing for Development, vol. 19. ACM (December 2010)
9. Bruff, D.O., Fisher, D.H., McEwen, K.E., Smith, B.E.: Wrapping a MOOC: Student Perceptions of an Experiment in Blended Learning. Journal of Online Learning & Teaching 9(2) (2013)
10. Milne, D., Witten, I.H.: An open-source toolkit for mining Wikipedia. Artificial Intelligence (2012)
11. Milne, D., Witten, I.H.: Learning to link with wikipedia. In: Proceedings of the 17th ACM Conference on Information and Knowledge Management, pp. 509–518. ACM (2008)
12. Wentzlaff, D.: Computer Architecture, `http://www.coursera.org/course/comparch`
13. Jurafsky, D., Manning, C.: Natural Language Processing, `http://www.coursera.org/course/nlp`
14. Boneh, D.: Cryptography 1, `http://www.coursera.org/course/crypto`
15. Wetherall, D., Krishnamurthy, A., Zahorjan, J.: Computer Networks, `http://www.coursera.org/course/comnetworks`
16. Ng, A.: Machine Learning, `http://www.coursera.org/course/ml`
17. Salton, G., Yang, C.S.: On the specification of term values in automatic indexing. Journal of Documentation 29(4), 351–372 (1973)