Video-Based Consensus Annotations for Learning: A Feasibility Study

Samuel Dodson
University of British Columbia. dodsons@mail.ubc.ca

Matthew Fong
University of British Columbia. mfong@ece.ubc.ca

Luanne Freund
University of British Columbia. luanne.freund@ubc.ca

Rick Kopak
University of British Columbia. r.kopak@ubc.ca

Dongwook Yoon
University of British Columbia. yoon@cs.ubc.ca

Sidney Fels
University of British Columbia. ssfels@ece.ubc.ca

ABSTRACT

Video-based learning is increasingly common in higher education; however, the video players available make limited use of logged interaction data to support and guide students' viewing. In this work in progress, we explore the feasibility of aggregating students' annotations of videos (e.g., highlights, notes, and tags) to identify "hot spots," which can signal areas of interest to subsequent learners. We conducted a deployment study with 315 undergraduate students using ViDeX, a video player designed for active viewing. We logged students' use of ViDeX for four months, and then aggregated and graphed their annotations across 13 instructional videos. Our results show that consensus annotations—the video content that has received attention from many students—may be a feasible, data-driven way to flag information for use by subsequent learners.

KEYWORDS

Video-Based Learning, Annotation, Interaction Data, User Interface Design, Information Visualization.

INTRODUCTION

Video-based learning is increasingly popular in higher education; however, most video players provide students with limited tools for learning. They commonly include only basic playback tools, such as play and pause, fast-forward, and rewind. This work is part of a larger project to design a video player, ViDeX (Fong et al., 2018), with tools for active viewing (Dodson et al., 2018). The current version of ViDeX supports user interaction with video, including highlighting, note taking, and tagging. This opens up the possibility of using this interaction data to enable students to learn from one another. We hypothesized that there would be sufficient agreement among students, such that aggregating and sharing data from their annotations could flag interesting or confusing parts of videos for subsequent learners. The objective of this work was to test this hypothesis and to develop a data-driven means of visualizing the aggregated interaction data for students.

Many types of interaction data can be aggregated, including the number of times an interval of content has been annotated. Marshall (1998) proposed that readers' annotations, such as their highlights and notes, could be aggregated in order to "identify n-way consensus, places in the text that all n of the readers… had agreed were important, or at least worthy of pulling out from the text." Like Marshall, we use the degree of annotation consensus as an indicator for agreement of points of interest. High degree of consensus among individuals may indicate what information they find confusing or interesting.

Video players are just beginning to make use of students' interaction data in order to support subsequent learners' video use. LectureScape (Kim et al., 2014) includes a visualization of aggregated view counts of a given video. By looking at the peaks and troughs of the aggregated view count plot, students can see which intervals of video content are most and least popular. LectureScape uses "byproducts" of students' navigation behaviors, which are implicit measures of students' video use; however, our data source is annotation, which is an explicit user interaction. MudSlide (Glassman et al., 2015) allows students to click on the visual content of a video to apply a "muddy point," a semi-translucent circle, to flag confusing content. Aggregated, explicit user interactions, such as muddy points, may provide more insight into what students find important or confusing than aggregated, implicit measures of video use, such as view counts. In this work in progress, we explore the use of aggregated summary plots and heat maps as visualizations for assessing degree of annotation consensus.

METHODS

We introduced ViDeX to 315 undergraduates in an introductory engineering course. Students were provided with 13 optional videos (8±3 minutes each) to watch before classes. The videos showed their instructor working through practice problems similar to those on their homework and exams. Students received a demo of the annotation features of ViDeX at the beginning of the semester, but they were not required to use the system. We logged students' use of ViDeX for four months, and then aggregated, normalized, and graphed summary plots and created heat maps of student's aggregated annotations for each video (Figures 1 and 2). Two videos were removed from analysis as they were annotated by fewer than 5 students. To smooth the plots of the logged data and enable comparison, the remaining 11 videos were binned into 10% time intervals, which ranged from 22 to 103 seconds depending on the video length. The annotations within each interval of a video were counted

and divided by the total number of annotators for the video. Consequently, multiple annotations by the same student were counted once. The normalized values for each interval are between 0 and 1, and these values were used to create the heat maps. The warmer colors represent higher degrees of consensus in highlights, notes, and tags and serve to identify "hot spots," or intervals of video content that have been annotated by at least half of students who annotated a video.

PRELIMINARY RESULTS

Of the 315 students in the course, 281 logged into ViDeX at least once during the semester. Of these, 77 students (27%) annotated, creating 593 highlights, 264 notes, and 1,178 tags in total. User interaction plots were created for each video (Figure 1), to provide visual summaries of annotations from all students. Highlights are displayed as horizontal bars. The color, length, and position of each bar matches the color, length, and position of the corresponding ViDeX highlights. Notes and tags are displayed as blue and red points, respectively. Figure 2 provides an example of a heat map for the same video, with the value in each interval representing the degree of consensus annotation. For example, a value of 0.5 means that half of those who annotated a video, did so in this interval. We found that all but one video had hot spots. The number of hot spots in each video varied between 1 and 3. Of all hot spots, 80% were in the first half of the video content. Hot spots, which may represent confusing or interesting video content, are now visible with the heat map.
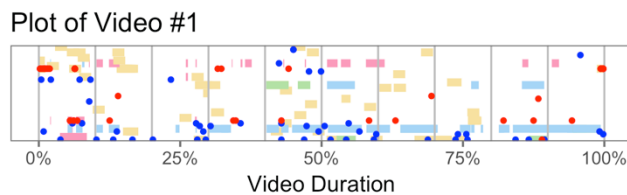

Figure 1: The summary plot of annotations for video #1. Each row represents a user. Highlights are displayed as colored horizontal bars, matching the color of the ViDeX highlights. Notes and tags plotted as blue and red points.
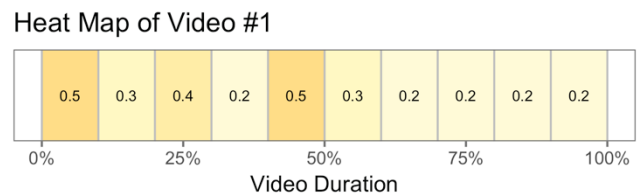
Figure 2: The heat map showing annotation (highlights, notes, and tags) consensus for video #1. The values and shades in each interval represent the degree of consensus.

CONCLUSION & FUTURE WORK

This initial study supports our hypothesis that there are discernable patterns of consensus in annotation across this set of videos, which could be informative to subsequent learners. We created two prototype visualizations, summary plots and heat maps, both of which could be displayed alongside a video timeline. We consider the heat map to be a good candidate for conveying degree of consensus annotation to users, as it is easy to interpret and scalable. Our next steps include: i) analyzing the aural, textual, and visual features of video in the hot spots to identify content that may trigger students' confusion or interest, and ii) evaluate the usability and value of aggregated annotations for students. Longer term questions include: does supplying this information change students' viewing behaviors, their level of engagement, or their learning outcomes?

ACKNOWLEDGEMENTS

REFERENCES

Dodson, S., Roll, I., Fong, M., Yoon, D., Harandi, N. M., & Fels S. (2018). An active viewing framework for video-based learning. *Proceedings of the Fifth ACM Conference on Learning at Scale*, 24:1–24:4.

Fong, M., Dodson, S., Roll, I., & Fels S. (2018). ViDeX: A platform for personalizing educational videos. *Proceedings of the 2018 ACM/IEEE Joint Conference on Digital Libraries*, 331–332.

Glassman, E. L., Kim, J., Monroy-Hernández, A., & Morris, M. R. (2015). Mudslide: A spatially anchored census of student confusion for online lecture videos. *Proceedings of the 33rd ACM Conference on Human Factors in Computing Systems*, 1555–1564.

Kim, J., Guo, P. J., Cai, C. J., Li, S. W. D., Gajos, K. Z., & Miller, R. C. (2014). Data-driven interaction techniques for improving navigation of educational videos. *Proceedings of the 27th ACM Symposium on UIST*, 563–572.

Marshall, C.C. (1998). Toward an ecology of hypertext annotation. *Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia*, 40–49.