Playback-centric visualizations of video usage using weighted interactions to guide where to watch in an educational context

Hyowon Lee*, Mingming Liu, Michael Scriney and
Alan F. Smeaton

Insight SFI Research Centre for Data Analytics, Dublin City University, Dublin, Ireland

The steady increase in use of online educational tools and services has led to a large amount of educational video materials made available for students to watch. Finding the right video content is usually supported by the overarching learning management system and its user interface that organizes various video items by course, categories and weeks, and makes them searchable. However, once a wanted video is found, students are often left without further guidance as to what parts in that video they should focus on. In this article, an additional timeline visualization to augment the conventional playback timeline is introduced which employs a novel playback weighting strategy in which the history of different video interactions generate scores based on the context of each playback. This includes whether the playback started after jumping forward or backward in the video, whether the playback was at a faster or slower speed, and whether the playback window was in focus on the student's screen or was in the background. The resultant scores are presented on the additional timeline, making it in effect a playback-centric usage graph nuanced by how each playback was executed. Students are informed by this visualization on the playback by their peers and can selectively watch those portions which the contour of the usage visualization suggests. The visualization was implemented as a fully-fledged web application and deployed in an undergraduate course at a university for two full semesters. A total of 270 students used the system throughout both semesters watching 52 videos, guided by visualizations on what to watch. Analysis of playback logs revealed that students selectively watched portions in videos corresponding to the most important portions of the videos as assessed by the instructor who created the videos. The characteristics of this method as a way of guiding students as to where to watch as well as a complementary tool for playback analysis, are discussed. Further insights into the potential values of this visualization and its underlying playback weighting strategy are also discussed.

1. Introduction

Increasing amounts of educational video materials are becoming available online. Coursera, one of the most popular Massive Open Online Course (MOOC) platforms, has over 3,000 online courses with 82 million registered learners as of March 2022 (https://about.coursera.org/), probably with millions of video materials of varying lengths available. When the COVID-19 pandemic hit in Spring 2020 forcing almost all universities to deliver their courses remotely, the amount of video content for students to remotely access further increased, with two main types of video content: (1) synchronous recorded lecture videos made available after the online, live lectures on Zoom or other video conferencing platform (typically lengthy, e.g., 30 min to 1 h or longer), and (2) video material prepared by the lecturer out of class and made available for students to watch asynchronous (typically shorter, e.g., 5–15 min, but can also be longer).

Regardless of the type of video material, online courses meant a considerable amount of videos for students to wade through and watch, often multiple video items per week, including lengthy lecture recordings as well as short videos to review. These videos typically become available on the university's Learning Management System (LMS) so that the videos for each week to watch are grouped, arranged, ordered, and possibly searchable for students to access: students are generally provided with some ways to find a suitable subset of videos that could satisfy their learning needs.

Once a video is selected and opened, however, students are usually left watching the contents without much support: (i) the inherent temporal nature of the video medium hides the full content from being browsed, and (ii) the typically unprofessionally- and minimally-edited, linearly organized characteristics of educational contents (mixtures of PowerPoint slides sometimes with animated annotations and/or talking heads, etc.) mean no clear or explicit structure to be followed within the videos. This results in either having to watch the full video even when not all of the contents are relevant/useful for the student, or blindly jumping back and forth hoping a useful point in the video will show up. Without instrumenting a lot of sophisticated within-video content browsing facilities (storyboards, keyframe-based browsing, structured shot/scene navigation, etc.) often the subject of experimentation within the video retrieval community (Lee et al., 2006), how could students be guided as to where to watch inside this type of videos?

Current analytic studies of video usage in the educational setting usually target *the instructors* as the beneficiary of the analytics, i.e., to show instructors how students are watching the videos, as there is a clear need for more objective insights and better visibility of students' viewing behavior to the teachers so that they can improve how they make and provide their video materials (Fong et al., 2019).

In this paper, we shift the focus of analytics *from the instructor to the students*, i.e., the usage of video contents watched by the student cohort can be analyzed and the results can be used to support students' navigation within the video content, an aspect which is missing in today's educational video platforms. By focusing on how potentially useful usage information could be fed back to the students so that it will guide them in choosing where to watch within each video, a *playback-centric visualization* was devised in which we developed a playback weighting strategy that takes into account various micro-level video interactions that students trigger and that are weighted differently. These weights are used to calculate an overall score for every 1 s window within a video, and we then present this as a video timeline.

By computing the cumulative values of individual video interactions (e.g., play/pause clicks or seek actions), and presenting the cumulative values of playback derived from differently weighted video interactions, the unique timeline visualization represents those parts within each video that have received the most interaction from students as it reflects how they watched when they were watching, e.g., were they skipping or perhaps using faster (1.25x, 1.5x, 2x) playback, as well as when during the semester they were watching the video. Our paper introduces this playback-centric visualization of video usage, in which:

- a playback weighting strategy is used and updated daily to calculate the second-by-second scores in each video based on judiciously chosen weightings for micro-level video interactions from students;
- a full-fledged system incorporating this strategy was built and deployed for two full semesters (2021 and 2022) at a university, and
- we analyse the usage data and the usefulness of this method, and share the insights gained from it.

The contributions of the paper are that by going through the above activities, we establish our understanding of the designed visualization, which can be expressed as our attempt at answering the following research question:

Does the proposed visualization of weighted playback by peers help students to selectively watch the more important parts within videos?

More analyses of the usage, perceptions and attitudes of students on an exclusively online learning environment during the pandemic (Chen et al., 2020; Coman et al., 2020; Popa et al., 2020) and attempts at capturing the students' attention level and monitoring it while they watch video or live lectures (Pham and Wang, 2015; Robal et al., 2018; Lee et al., 2021; Srivastava et al., 2021) are all very important, and the support for better guidance

for students in more selectively/efficiently reviewing the video contents is a notable gap identified and addressed in this paper.

2. Related work

There is a large body of literature in learning analytics for online platforms in general, e.g., comprehensive reviews including (Ihantola et al., 2015; Dawson et al., 2019), and a moderate amount of literature on analysing video watching behavior in an educational context. When it comes to detailed analysis looking into the within-video-level watching behavior such as the proportions actually watched and parts of videos that had more interactions, the number of publications decreases to a handful. The increased adoption of MOOC courses in universities certainly increased an interest in deeper understanding of the issues of online learning (e.g., decreased motivation, delayed feedback, feeling of isolation, etc.; Yusuf and Al-Banawi, 2013; Coman et al., 2020) and of the analysis of video material usage and student behavior, the bulk of the latter will be reviewed in this section. We expect a more recent line of such studies to appear soon as the COVID-19 restrictions in universities resulted in more intense e-learning thus an increased amount of video-based online classes.

2.1. Within-video usage analytics in educational context

Video usage analytics outside the educational context have been many and diverse including market surveys on why people re-watch the same commercial video content (Bentley and Murray, 2016) and market studies using the commonly used "views" and "likes" counts in YouTube. The use of "views" in particular, is the most basic, dominant metric used in learning analytics studies where the main focus is the usage of the whole learning management system (LMS) in which the video usage is a part (e.g., Jovanović et al., 2017).

Other than views, there are other metrics typically used for video analytics including "play rate" (ratio of playing a video to the number of links usually in the form of a thumbnail frame image, shown to the viewer) and "view length" (how long the viewer watched the video). In the context of educational videos and pre-class asynchronous learning materials, "punctuality" is sometimes used to measure whether the students watched the video before the class time. Common metrics used to look into within-video watching behavior include "coverage" (ratio of playback portion to the whole duration of the video), "peak" (most frequently interacted part of the video), and "drop-off point" (where the viewer stopped watching in the video). In the video usage analytics in the educational domain, some combinations of these metrics have been used.

An in-depth interview with 16 post-secondary instructors who use videos in their teaching (Fong et al., 2019) showed that a more in-depth analysis of video watching behavior of students would be very useful for them to assess the video quality and how they could revise accordingly. It is apparent that lecturers will benefit if there are more detailed information available to them on how video materials are viewed by students.

On the one hand, the majority of the literature on video playback interaction analysis to support such insights in an educational context comes from around 2013 when large-scale MOOC courses became popular and a greater variety of in-video interaction types (pause, seek, speed change, etc.) became the norm on these platforms. This meant that the analysis of these click-level interactions became technically feasible in practice. On the other hand, the degree of detail that the studies look at within-video interaction varies, from simple view counting presented over time to second-by-second cumulative calculation by different interaction types.

For example, one study (Dazo et al., 2016) looked at the usage of 25 video materials across three consecutive semesters in a flipped classroom setting. Using views, coverage and punctuality as the main metrics, they found that students' academic performance was connected to some of the video viewing behavior, e.g., previewing the videos before the class. Looking at the watching behavior inside the videos was through the metric of coverage, generally proportional to the view counts. Another study found out two distinctive clusters of students based on their video watching behavior in a 15-week course on Matlab programming (Moore et al., 2021), where the high-performing cluster of students also showed significantly higher coverage in their watching, compared to the low-performing cluster.

Moving beyond the overall view count and coverage, an analysis of a MOOC course involving 48 videos (Atapattu and Falkner, 2017) counted the number of video interactions (play, pause, speed change, etc.) that happened at each second in each of the videos, presenting the counts on a timeline within the video. While their focus was to use video transcripts to analyze the discourse and correlate with those video interaction counts on the video timeline, the counting of interactions at each second of the videos and seeing the highs and lows of them on a timeline of each video is a promising approach to reveal the minute details of interactions as they happened while the videos were playing.

Going even deeper into within-video interactivity, Li et al. (2015a,b) related the frequencies of different types of within-video interactions (pause, skip, replay, etc.) to the understandability of video contents (i.e., difficulty of learning). Of particular interest is the metric "replayed video length" that summed all played video durations by seconds that were re-watched (concluding that less re-watching indicated higher video difficulty). Differentiating different types of interactivity is certainly necessary if we want to understand the playback

usage in more detail, and we describe later how in our study the different types of interactions will be assigned to different weightings in cumulative scoring inside each video.

A study in Kim et al. (2014a) used similar second-by-second bin counting where play and pause interactions increased the count thus forming the "peaks" in the usage timeline. They also used simple image similarity comparison between frames from every 1 s in the video to identify the points where high differences appear and thus help categorize different types of transitions in video presentation (e.g., from PowerPoint slides to talking head). They show a prototype user-interface where different types of interactions (views, replay, skip, etc.) can be selected to see its cumulative peaks within the video.

Further related studies and analyses include video analytics that used aggregated time series to present interactivity within the videos. This was used to check those parts of video contents in a specially designed 10-week course (Giannakos et al., 2015), correlating the video playback time with student engagement level in 4 edX courses that happened in the same semester (Guo et al., 2014). This took a cognitive psychology angle where grouping of different sequences of video interactions into higher-level "behavioral actions" to help predict the students' usage behavior (Sinha et al., 2014).

Table 1 shows a summary of seven video usage analysis studies (2014–2021) some of which addressed the above, selected as good representative studies in the literature that we reviewed. It can be observed from the table that:

- Educational video contents span from a few dozen to a few hundreds of videos, and the length of each video varies from a few minutes to 45 min;
- Various sets of metrics are used in different studies;
- There are attempts at correlating the watching behavior to some overall score (e.g., grades);
- Temporal mapping within videos (the last three rows in the table), attempt at within-video locating of usage behavior, could be explored more.

The studies that apply the last point (temporal mapping within video) count video interactions at second-level points inside each video, including clicking of play/pause buttons, replay, etc. This allows them to show the cumulative effects of usage as the peaks of these interactions along the timeline. The peaks then point to the video contents, at those points.

It is this temporal mapping of within-video usage that our paper builds upon, where the values to be presented are not the counting of interactions but the accumulated playback scores that take into account various video interactions, each weighted differently, in order to reflect the way students actually watched the videos. We expect this will better guide students as to where to watch.

2.2. Within-video recommendation and timeline visualization

Typical usage of video content has been to linearly watch from start to finish, whether short (e.g., a 5 min YouTube clips) or long (e.g., a 2 h movie), thus not requiring any elaborate visualization scheme to help to navigate or jump within the videos other than indicating where the current point of playback is in time. Educational video content, on the other hand, will generally benefit from features which support within-video navigation/guidance as the videos in educational domains typically have a structure, a progression and different types of presentation, and students might be looking for certain parts within them. There are cases where more use of video control features such as skip/pause was negatively associated with students' performance (e.g., Joordens et al., 2009), though that work pre-dates the huge uptake in online learning over the past decade and in particular as a result of COVID. Thus, thinking that providing various interactive video playback features will be beneficial for students may be only an assumption. Different learning strategies used by students (e.g., simply memorizing vs. deeply understanding) along with the topic types of the course will have impact in deciding how beneficial the video control features are (Le et al., 2010).

Some of the video analysis studies in educational contexts, including some of the earlier-mentioned ones, do feature visual guidance on the playback timeline. For example, TrACE UI shows the portions of video that they already played (Dazo et al., 2016) on its video timeline, guiding students on where to play or avoid playing; Video timeline shows peaks and highlights that are dynamically changing as the usage increases (Gajos et al., 2014). YouTube's video analytics page features a "key moments for audience retention" section in which the usage frequency is presented on a timeline (duration) of a video, visually showing the peaks and troughs presented based on playback as well as other interactions such as sharing with others at those points. To support video-based learning of how to use Photoshop effects, the video timeline shows the highlight points for step-by-step navigation (Kim et al., 2014b).

Although not particularly for an educational context, Video timelines with bookmarks allows smooth pre-fetching and the indication of buffering has been proposed (Carlier et al., 2015); Ice hockey goal events are highlighted and shown as mini icons on the timeline of the playback UI (Müller et al., 2010). Many experimental video retrieval systems in the field of multimedia (e.g., the annual Video Browser Showdown showcase; Schoeffmann, 2019) have been experimenting with more sophisticated and advanced within-video highlights where the retrieval unit is sub-video parts (e.g., a frame, camera shot, scene, etc.) and the searcher is supported by various visual strategies to highlight within-video match points within long videos. Examples include highlighting the portions of

TABLE 1 Comparison table amongst seven representative analysis studies of educational video usage (this table shows only those metrics used and applied as the main focus of the investigation in each study).

	Sinha et al. (2014)	Kim et al. (2014a)	Giannakos et al. (2015)	Li et al. (2015a)	Dazo et al. (2016)	Atapattu and Falkner (2017)	Moore et al. (2021)
Course	Coursera (prgm.)	4 edX courses (stat., AI, chem., prgm. in 1 semester	Specially designed course (7 weeks)	2 Coursera courses (prgm. and sig. proc.)	1 course (prgm.) revised for three consecutive semesters (15 weeks each)	AdelaideX (edX) over 6 weeks	Prgm. course for two semesters (15 weeks each)
Class size	22 k students	128 k students	11 students	32 k students	23, 19, 24 students each semester	26k registered students	215, 275 students each semester
Course topic	Functional programming using Scala	Intro. CS, programming, statistics, AI & Solid State Chemistry	Use of information for problem-solving, research & decision-making	Reactive programming & digital signal processing	Intro. Java programming concepts	Intro. Processing-JS programming concepts	Intro. MATLAB programming
Archive size	48 video lectures	862 videos (average 7 min)	7 videos, for pre-class preparation only	94 videos, total 300k sessions (visits)	25 videos for pre-class watching, avg 45 min	48 videos, average 3.63 min	70 videos, each 5-10 min
Video interaction used (metrics)							
- View		✓	✓		✓		✓
- Coverage		✓			✓		✓
- Punctuality					✓		✓
- Counting all clicks						✓	
- Play/pause	✓	✓		✓		✓	
- Skip/Replay	✓	✓	✓	✓		✓	
- Speed change	✓			✓		✓	
- Exit/quit		✓					
- Caption on/off						✓	
- Play duration	✓						
Differentiating among interaction	✓	✓		✓			
Clustering students/behavior	✓						✓
Correlating the interaction to a grade/outcome	Dropout		Video assessment scores	Perceived difficulty	Course grades		Test scores each week
Temporal mapping within video							
- Interaction counting		✓	✓			✓	
- Visual shifts in video		✓					
- Discourse						✓	

a video timeline for a particular type of scene (dialogue, action, or montage) in movies (Lehane et al., 2006), indicating the boundaries between short video chunks in a longer, concatenated video (Lee et al., 2017), and various ways of highlighting retrieval results on the video timeline (Gaughan et al., 2003; Cooke et al., 2004; Leibetseder et al., 2018; Sauter et al., 2020). We do need to see yet more effort in applying these sophisticated within-video timeline visualizations in an

educational context. Judicious applications of some of these features could help students more efficiently learn the contents by more selectively picking and playing the parts of videos that are useful for them.

3. System description

The system we built and deployed is a web-based archive of educational video materials made available throughout a semester, used in conjunction with the university-wide Learning Management System. The video materials, once they are available on the system for watching, can be played with conventional playback-related functions. While all user interactions throughout the web interface are logged, the system's visual guidance feature which this paper introduces and the analysis used to drive its operation, uses mainly the playback-related interaction portions of the logs such as play, pause, and speed change as well as indirect, contextual usage logs such as whether the playback is happening as the window of focus on the student's screen.

Figure 1 shows how the system works. The lecturer uploads video clips that s/he created and offers the students to watch as part of the course material (1); the system will use a specially formulated playback weighting strategy to calculate what the score for each of the second-by-second bins in the timeline should be (2); scoring adds a time-decay function to construct the playback visualization for students (3); the resultant UI is then presented to students (4); finally, students interact with the system by playing, pausing, changing speed, etc. and this interaction is fed back to the system to re-calculate the scores using the weighting strategy (5). The cycle (2)-(3)-(4)-(5) continues through the semester as the shape of the timeline visualization evolves over time.

In the remainder of this section, we describe the user interaction and visualization, followed by the mechanism used to generate it.

3.1. User interaction and the usage visualization

The web interface allows students to log in by using their student email address and a unique eight-character code for the course. Once logged in, the student is presented with a list of available video materials on the left side of the screen (see Figure 2). Selecting one will load the video playback panel and the associated usage visualization above it.

The visualization feature is a time-based graph, aligned with the playback timeline of video content (panel above the playback screen in Figure 2), where the height of the yellow shape suggests video usage to students so that they could watch those portions within videos that show high levels of usage by their peers.

Using the yellow contour of the graph as a hint to select parts of the video to watch, students engage in selective playback of the video, data from which will, in turn, be taken into account using the playback weighting strategy to contribute to the shaping of the contour for other students. Establishing the usefulness of this visualization in helping students watch the important parts of videos is the focus of the analysis presented later.
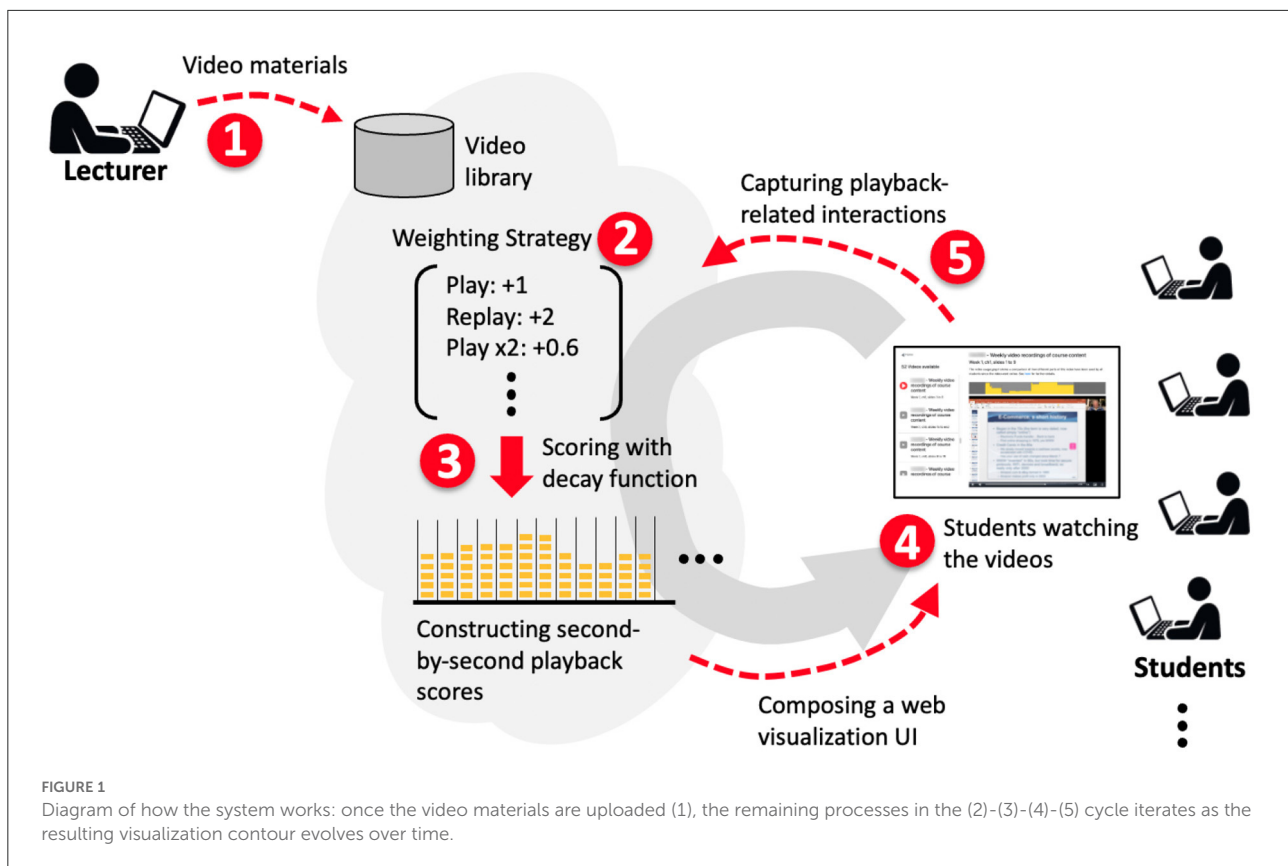
3.2. Playback weighting strategy for the visualization

The height of the yellow areas within the visualization is the result of our playback weighting strategy, the core technical component of the system.

Each video is divided into 1 s windows and each window starts with an initial score of 0 which is incremented every time any student plays or skips it. Every time any part of the video is played, that part will gain a score increment, thus over time as the usage increases the scores in each second-by-second window on the timeline will increase.

Since there are various playback-related interactions a user can perform at any time during video playback, the system uses the type of interaction that happened to calculate the score increment for that second: thus the playback weighting strategy determines the shape of the visualization, eventually serving as the guidance to students on what parts of the video to watch. The strategy is as follows:

- **Play**: As the most basic scoring strategy, the window gets **+1** when that portion is played. This is similar to calculating "peak profile" in which play events are counted to show the peaks (Kim et al., 2014a). However, if the video playback window was not the window of focus on the student's screen when the student was playing it, then this increment is **+0.25** only, bringing down the score for cases where the student may be reading email or doing something else while half-listening to the video.
- **Replay (Seek backward, i.e., within the same session)**: All 1 s windows between the previous playback point and the point it skipped back to will each gain **+2** score, so each gets a cumulative +3: +1 from the initial playback and +2 from the replay). Re-watching could mean more difficulty in understanding that part of the video (Li et al., 2015a), although there could be other factors, e.g., returning to missed content, following the steps in tutorial-type videos, hearing again non-visual explanations, etc. (Kim et al., 2014a). Whatever the reasons, the parts which students replayed are parts to watch out for and point out during students' course review.
- **Faster play:** Playback at 2× (double speed) will gain **+0.6** and if the window is not the window of focus then it will

FIGURE 1
Diagram of how the system works: once the video materials are uploaded (1), the remaining processes in the (2)-(3)-(4)-(5) cycle iterates as the resulting visualization contour evolves over time.
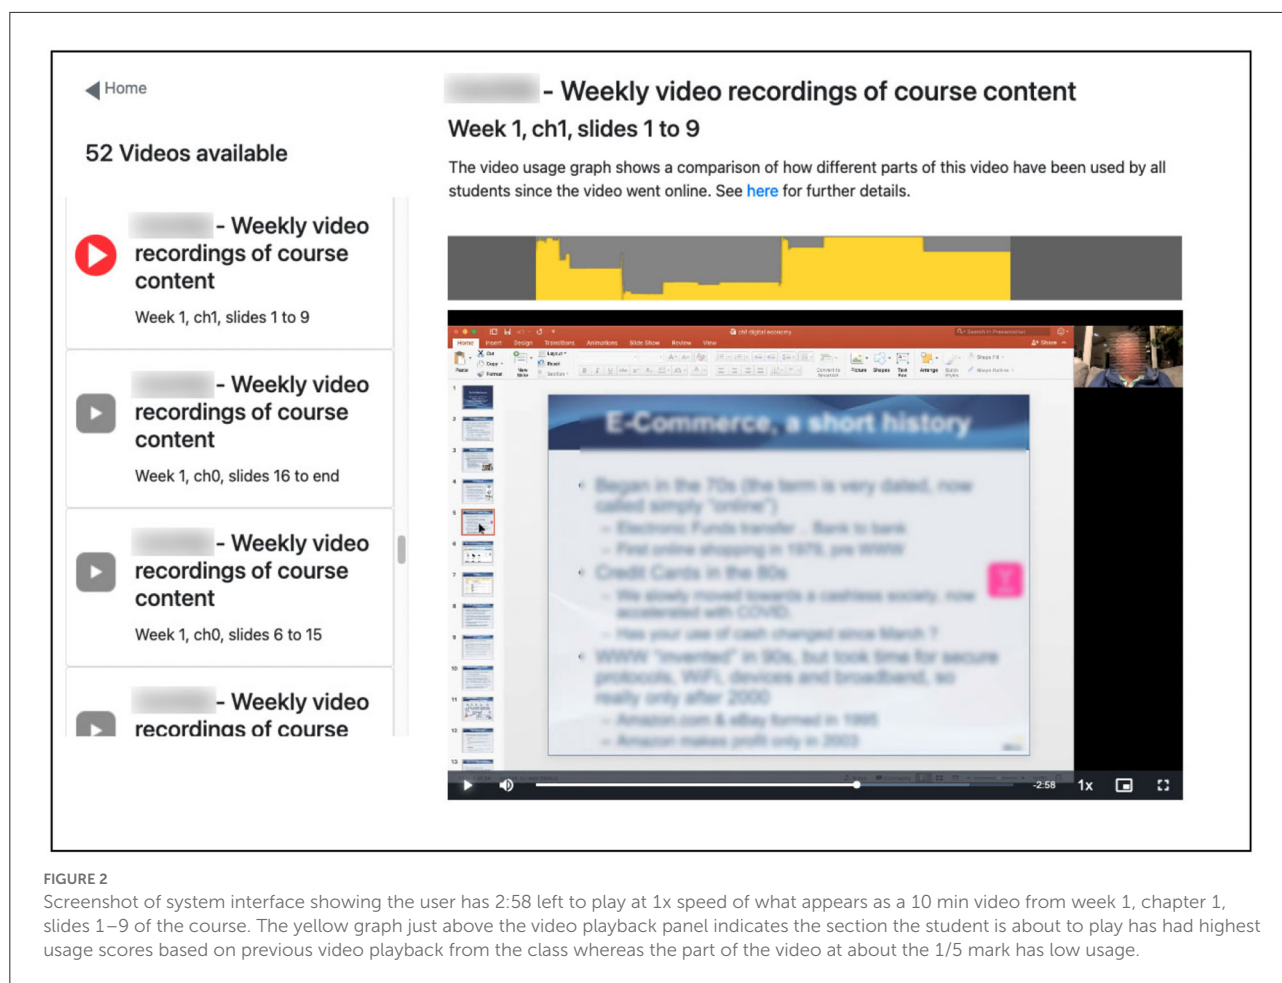
be **+0.2** only because when the student is attending to another window, the double-speed playback is too fast to properly comprehend. Playback at $1.5\times$ will gain **+1.5**, and **+0.5** if the playback window is not the window of focus. In this study we assume that playing at $1.25\times$ speed will gain the same as playing at $1\times$ speed for student learning. Recent studies consistently show that playing educational videos at $1.25\times$ resulted in better learning outcomes than when played at normal speed (Lang et al., 2020; Mo et al., 2022), and the benefits and drawbacks of other speeds (e.g., $0.75\times$ and $1.75\times$) are worth further investigation in order to incorporate them into our weighting strategy.

- **Skip (Seek forward)**: There can be different reasons from a cognitive processing point of view as to why students want to skip familiar parts of a video or continue watching even if already familiar (Sinha et al., 2014), but skipping typically implies that the student finds those portions easier to understand (Li et al., 2015a). If a student skips forward from the current position at minute S0 by a segment of video then windows within 1-, 2-, and 3-min segments directly following the segment S0 will get score adjustments as follows: S60: **−0.3**, S120: **−0.2**, and S180: **−0.1**. The rationale for this deduction in scores is that the student must have had an idea what was coming up next, after point S0, but anticipated it as being not interesting or useful for

her/him at that point in time and thus less likely for other students to also find it interesting or useful.

3.2.1. Adjusting the score over time

Overarching the playback weighting strategy is a **decay function** based on recency of playback, with the most recent days of playback interaction being more meaningful or useful than prior to that. The scores calculated by the above strategy are re-calculated from the interaction log each midnight. In this way the score increments (both + and −) as above are called our BaseIncrements = (+1.0, +0.25, +2, +0.6, +0.2, +1.5, +0.5, −0.3, −0.2, −0.1) and the system makes those the actual increments on day 0. Then on day 1 it makes those increments each multiplied by 1.1 before adding, on day 2 the base increments multiplied by 1.2, and so on. The effect is that on day 10 we have a score for each 1 s window which has a 10-day linear decay function so that something played on day 9 has twice the value of something played on day 0 and by day 20 we have a score which has a 20-day linear decay with the half-life being 10 days. This continues indefinitely. The decay function is useful for the lecturer who wants to understand the factual usage statistics of each video segment, but it serves an important role that will influence the shape of the usage visualization for students who need to be selective at the time of reviewing the video.

FIGURE 2
Screenshot of system interface showing the user has 2:58 left to play at 1x speed of what appears as a 10 min video from week 1, chapter 1, slides 1−9 of the course. The yellow graph just above the video playback panel indicates the section the student is about to play has had highest usage scores based on previous video playback from the class whereas the part of the video at about the 1/5 mark has low usage.

The scores for the 1 s windows are normalized within each video usage graph at display time making the visualization less susceptible over time to any maligned attempt at artificially inflating scores by jumping to or repeatedly playing an obscure segment of video.

A study found that the points in videos with more video interactions indicate those parts requiring higher-order thinking and more cognitive skills by students (Giannakos et al., 2015). Instead of scoring highly those parts with more interactions regardless of the type of interactions, our weighting strategy attempts a more fine-grained differentiation amongst different video interactions and uses the cumulative scores from it as overall guidance to offer to students.

4. Course outline, deployment, and data captured

The course for which the system was deployed was an undergraduate-level course on data analytics for marketing applications delivered in the Spring semester (18 January–15 May) in 2021, and again the following Spring semester (10 January–26 April) in 2022.

Since the Spring semester in 2021 was a lockdown period due to the COVID-19 pandemic, the course was delivered fully online: Zoom lectures were conducted live while video recordings of them were made available for students for later asynchronous review as they wish. The University's virtual learning environment was also used to present other course material including links to YouTube clips, news articles and other websites as well as the PowerPoint slides used for the online lectures and for creating the video materials and worked examples of data analytics. As an additional source of course material, the Professor prepared a set of 52 short form video clips for students to watch in their own time in addition to various links to the other materials. In each week, the materials suitable for that week's lecture topic became available as entries on the course's LMS site. The number of students enrolled to the course in the first deployment was 131.

The Spring semester in 2022 was conducted with regular face-to-face in-class lectures as the main mode of course delivery and those lectures were also streamed live and Zoom

TABLE 2 Topic of each video.

v1	Course overview	v27	Trusting data visualizations
v2	Course content, exams, assignments	v28	Data visualizations reveal data
v3	Online resources, past student feedback	v29	Sample classic data visualizations
v4	Our digital economy	v30	Why data visualization is important
v5	ECommerce, how and why it works	v31	Perceive, interpret, comprehend. Pareidolia
v6	Traditional marketing	v32	Worked examples of human visual processing
v7	Digital marketing and analytics and "free"	v33	Chart types, bar charts, connected dots
v8	What are our digital footprints	v34	Radar, polar, range, box and whisker, histograms
v9	GDPR introduction	v35	Treemaps, sunburst, bubble, sankey, line charts
v10	Tracking web browsing activities	v36	Choosing charts to use, single numbers
v11	GDPR Principles	v37	Correlations, regression, pie charts
v12	GDPR details and how people differ	v38	Overview of color, RGB and HSL
v13	Digital footprints and personality profiling	v39	Nominal data, ordinal data and color
v14	Descriptive statistics, probability distributions	v40	Data visualization and color
v15	Normal distribution, standard deviations	v41	Animation in data visualization, Tableau
v16	Poisson distributions, scatter plots, correlations	v42	Logistic regression
v17	Sampling bias, significance tests	v43	Machine learning overview
v18	Chi test, introduction to regression	v44	Performance of machine learning models
v19	Linear/logistic regression, regression principles	v45	Market segmentation, Google and Facebook
v20	Sample regression	v46	Algorithms for market segmentation
v21	Linear regression in Excel	v47	Using transaction, interaction & external data
v22	Data in normal form, data imputation	v48	Where to find external open data
v23	Joining or merging data, VLOOKUP in Excel	v49	Data analytics in politics
v24	Data analytics pipeline/visualization overview	v50	Digital footprints to infer personality
v25	Storytelling in data analytics	v51	Facebook, Cambridge Analytica & US election/Brexit
v26	Two data analytics case studies	v52	Living trace-free on the internet

recordings were made available, further augmented with the same additional online material as in the previous semester including the same set of video clips prepared by the Professor from 2021. Each of the video clips was pointed out and recommended for students to watch as the course progressed through the face-to-face lectures. The number of students enrolled in the second deployment was 139.

The video materials were screencasts created by the Professor who taught the course. Each video was about 10 min duration and 3–5 such videos were added for each week throughout the semester. The topic for each of the videos is shown in Table 2. The time of the addition of each video depended on factors including the availability of certain materials to be mentioned in the video and time constraints of the Professor, resulting in some materials on later topics to be available earlier in the semester, and vice versa. A total of 52 videos had been made available to students in both deployments, and the usage of these 52 videos facilitated by the visualization during the two deployments is our main focus in this paper.

For the duration of each of the two semesters, enrolled students used the system actively playing and re-playing video content as part of their studies thus feeding into the playback

usage analysis which, in turn, helped shape the contours on the timelines for each video that guided them where to watch.

A manual annotation of the most important part within each of the videos was carried out in order to check whether the visualizations actually helped students locate the most important parts of videos in their study. The Professor who developed the video materials and delivered the course went through each video, marking the start and end points within each video that were judged as the most important part of that video, not knowing the actual playback usage that happened during the semester. The only constraint imposed on the Professor's annotation was that for each video there was to be only one continuous part as the most important, that is, not multiple different time blocks. The duration of this part was up to the Professor to decide purely based on the learning content. Although there might be multiple points that a lecturer sees as important in courses where video materials are longer in duration (e.g., 30 min or 1 h or more), the videos in this course were short and the Professor specifically designed each video to be about one particular concept. Since the course and its video materials were all developed by the same Professor, there was no need to employ multiple assessors to indicate

TABLE 3　Overall playback usage during the two deployments.

Number of students	270
Total number of videos	52 (average duration: 10 m 08 s)
Total number of sessions*	1,521
Total hours of playback	130
Average minutes of playback per student	28.9

*A session is defined as a student watching at least one video without being inactive for more than 10 min.

these important parts and to then check cross-annotation for consistency.

It is worth noting that the annotation of the important parts within each video is inevitably a manual task that had to be performed by the Professor in order to answer the research question in our study, and thus for the regular deployment of the system such a manual effort is not required.

5. Findings and analysis

Overall usage of the video materials using our system is summarized in Table 3, showing combined usage from both semesters. The analysis and our evaluation of the usefulness of the proposed playback weighting strategy and visualization uses the combined usage data collected from both semesters. Since the video clips were supplementary to the main lectures by the Professor either by Zoom or face to face or both, as well as to other online materials, students were not required to watch all available videos (though they were encouraged to), which explains the overall low playback usage (an average 28.9 min of playback per student). This indicates that when students did use these videos it was not their main source of course material and could have been used to clarify some complex aspect of the course, though this is speculative.

5.1. Playback usage according to playback weighting

In order to answer our research question, we examine the playback usage according to the timeline visualization shaped by our playback weighting strategy. Figure 3 shows the playback behavior within each of the 52 videos according to our weighting strategy but without reflecting the decay function (designed to factor in the date at which the student is watching the video). Each graph represents the second-by-second playback scores presented for each video, according to the playback weighting strategy described in Section 3.2.

The blue contour lines in each video graph show the playback usage of each video. As can be seen, the contours are overall quite different-looking than those that present the counts of video interactions since here we present the cumulative,

played parts of videos with the influence of the playback weighting strategy: there are fewer sharp peaks but a greater number of different levels of plateaus where playback continued for some time.

To begin with, an initial observation from Figure 3 is that all graphs start with high playback usage. Students start watching a video by clicking on the play button, and as the starting point is always at the start of the video when the video is first loaded, the playback score is always counted at the very start of every video. But after playback starts, almost always the students will skip into the middle of the videos, making the playback scores at the start immediately drop. Exceptions to this behavior can be found in only a few videos, for example, the video labeled v45 in Figure 3 shows that from the very start, playback usage was well maintained for most of the video; the video labeled v31 shows a start with a low usage but steadily increasing until the end of the video at around the 5 min point.

The timeline visualization that students actually saw and were guided in choosing the portions to watch are the decay function-enhanced versions of the Figure 3, that showed slightly different contours depending on when they watched the videos. It is not possible to accurately estimate how the usage might have been without the decay function influencing the visualization as the usage evolved over time with the function, but by back-tracing the contours during the semester, we can estimate its effect (see Section 5.2).

5.1.1. Effectiveness of the visualization: Highlighting the important parts

Checking the playback usage together with the important parts manually annotated by the Professor is a central task to answer our main research question. In Figure 3, the green areas in each video graph are the important part in each video marked by the Professor who gave the course and annotated this independent of the actual usage. The marking range for each video was then compared to the playback usage visualization for each video, to check whether the parts with high plateau of playback scores (i.e., the part the students were guided to and watched) is where the Professor marked as the most important parts. The green check symbol on the right side of each video graph in Figure 3 shows that the highest playback regions matched the important parts of the videos, while the red "X" symbol shows that they did not match. Out of 52 videos, 38 matched (73.1%): that is, 38 out of 52 playback visualizations correctly guided students to watch the most important parts within the videos and they actually watched those portions.

This deserves more detailed analysis. For example, video v6 (see Figure 3) shows a low usage for the first 5 min or so, then suddenly peaks maintaining the plateau for 30 s before flattening, and the peak point coincides with the important part as judged by the Professor. It is encouraging to see that the natural playback usage reinforced by the visualization, results in
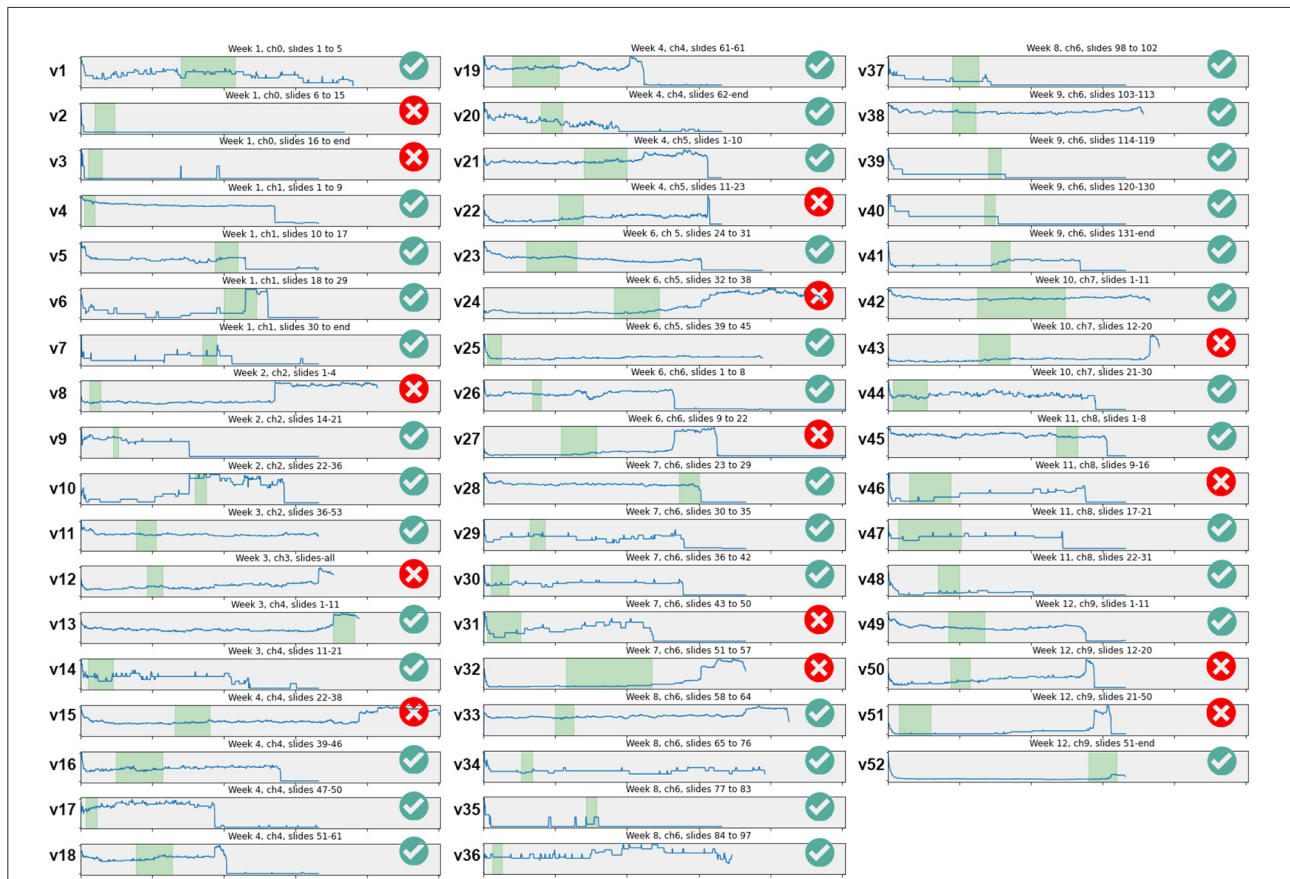
FIGURE 3

Within-video playback usage for each of the 52 videos using the playback weighting strategy (without decay function), each video labeled v1−v52. In each graph, x-axis is video duration, y-axis is a normalized overall playback score for each second. In each video, highlighted in light green is the most important part in the video; ✓ or × symbols indicate whether the highest plateau covers the most important part (i.e., students watched the important part) or not.

students watching that important part of the video. As another example, v41 shows relatively low usage throughout, but halfway into the video at around 3 min point, the line slightly goes up and remains plateaued: the start of that plateau is the important part and students correctly catch that part.
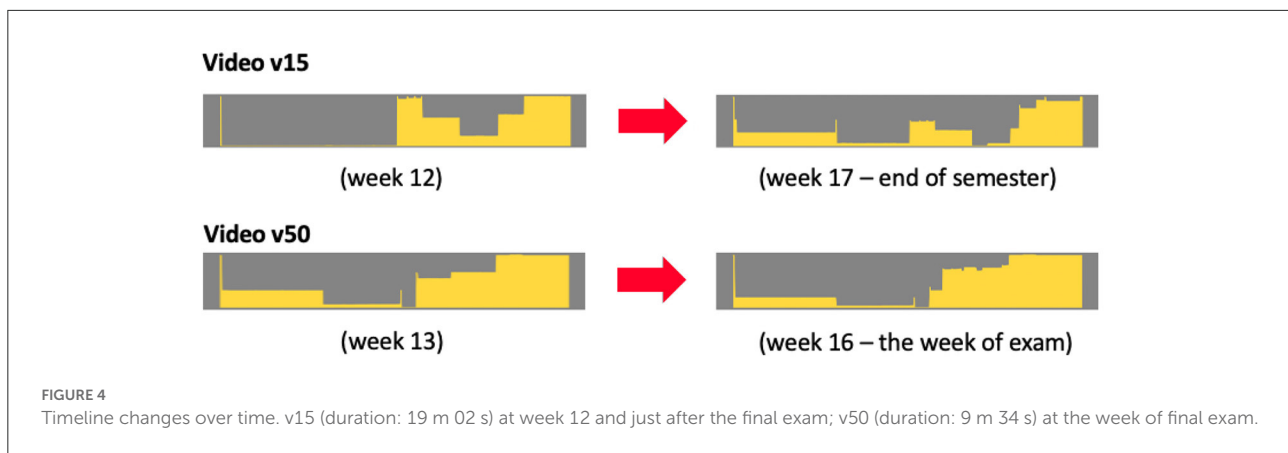
On the other hand, v42 shows a middle third of the duration as the important part, yet the actual playback was pretty much a high plateau all throughout the duration of the video. The visualization ends up not differentiating any part and is thus unable to effectively guide students to be more selective, yet most students who opened this video did watch the important part.

5.2. Contour changes over time

As students play the videos, the timeline visualization for each video changes as it is updated daily. Figure 4 shows two example timelines captured near the end of lecture weeks

(Weeks 12 and 13 respectively) and again at/after the final exam, during the 1st deployment (since the visualizations that the students actually saw cannot be based on aggregated data across two semesters). In case of v15, the topic covers the normal distribution in a set of numeric values and the concept of how the standard captures the spread of data values. The parts played most are the second half by the end of lecture classes as students grapple with the concept of what is a standard deviation. Approaching the final exam the playback portions changed and the week after the exam the timeline shows more coverage across the video, ending up with two moderate plateaus at the start and middle followed by the highest region nearer to the end of the video.

On the other hand, v50 which is about how our digital footprints can be used to predict our personality using the 5-factor personality model, this seems to maintain the overall playback usage after the high use at the exam, though slightly lower at the first 3 min or so, and higher at the latter 4 min.

FIGURE 4
Timeline changes over time. v15 (duration: 19 m 02 s) at week 12 and just after the final exam; v50 (duration: 9 m 34 s) at the week of final exam.

The change of contours becomes more noticeable as the usage becomes higher, and in our system, this understandably occurred right before the final exam period started. Whatever contours shaped thus far as a result of relatively low but continuous and gradual usage during the main part of the course are then picked up as the hints to where to watch when the urgency of learning increases at the impending final exam (the danger of "bandwagon effect" and how our algorithm is resistant to it is discussed in the next section). However, the exact pattern in this more noticeable contour change near the exam time is not consistent and dependent on the video contents, as seen, for example, the way the peaks near the end becoming more spread out throughout v15 as the result of more usage due to final exam, and maintaining the peak near the end of video and becoming increasingly so due to final exam in v50, clearly contrasting to v15.

6. Discussion

6.1. Playback-centric view as complementary to video interaction view

The contours of the within-video graphs are quite different from those that present the counts of video interactions (e.g., Kim et al., 2014a; Giannakos et al., 2015; Atapattu and Falkner, 2017). Influenced by the playback weighting strategy that incorporates video interactions but presented not by those interaction points but by the played parts, the playback-centric presentation is more range-based (i.e., from time A to time B) than point-based (i.e., this point in time). The message the visualization gives is the suggestion of the parts of the video to watch and is thus slightly different from the presentation of where the interactions occurred most often. The playback duration-based usage on the timeline is an additional aspect of the usage which is complementary to the interaction point-based usage, adding further dimensions to understanding the details of the use of video playback.

6.2. High plateau = important part?

In this study, we attempted to assess the usefulness of the visualization (and its weighting strategy) with the broad assumption that the system's visualization is effective when the most-watched parts in videos correspond to the most important parts as determined by the Professor. This assumption may be naive since the parts where the students need to review more frequently or need to think more deeply might not be same as the most important parts identified by the Professor. Perhaps the more difficult parts, or the parts that omitted examples or were explained poorly, might have been watched more and thus the more difficult or poorly-explained parts are not necessarily the most important parts. However, as an initial deployment of the visualization, such an assumption could still be useful without complicating or adding further intervention steps during and after the deployment. This needs to be kept in mind in the interpretation of our findings.

Currently our playback weighting strategy is purely based on playback usage and influenced by the time-decay function. For 38 out of 52 videos in our study, the parts the students watched the most were largely the parts that the Professor judged as most important, but for the remaining videos this was not the case. The weighting strategy could incorporate content-related factors, e.g., the instructor's indication of important parts in each video, perhaps by ranking a few different parts of the videos, to influence the baseline score of each second-by-second bins. However, the added manual input required from the course instructor would be an issue especially as the number of courses and videos increases.

The timeline visualization changes over time as students continue to play different parts of different videos. Students will not benefit from the visualization in the early stages of the semester as not many parts will have been played and not a lot of noticeable contours will be shown. They benefit as the timeline plateaus become more prominent, only after it becomes clear where peaks are and this comes only after playback usage has occurred.

With this in mind, we now ask whether the proposed visualization helps students selectively watch more important parts within the videos and the answer is reasonably so (73.1%), considering no explicit information or intervention was offered to influence video usage. Furthermore, given our approach of relying on usage data to feed into the system, the visualization is likely to improve its guidance to important parts if there was a pedagogical mechanism to increase the overall usage by students, e.g., enforcing the watching of the videos as mandatory, rather than as supplementary material as in this course.

What happens when the high plateaus are not corresponding to important parts of the video? If so, it could mis-guide the students and potentially fall into a negative bandwagon effect as happens in social networks (Wang et al., 2015), in recommender systems (Sundar et al., 2008) and perhaps take advantage of the effect in e-commerce and IoT products (Choi et al., 2015). Unintended mis-guiding formation of the usage visualization is less likely to occur as the decay function embedded in the playback weighting strategy makes the shaping of the contours somewhat resistant to occasional random watching on obscure points within the videos. Initial "seeding" of important parts in each video may be a way to further reduce such negative effects.

Finally, the videos watched by the students in this study were the supplementary learning material in addition to the formal online/face-to-face lectures as well as other website links and YouTube resources that were made available for students. The playback usage pattern within the videos (thus the shapes of the contours) is expected to be different if the videos are to become the core delivery mechanism as is the case in many studies referenced in this paper so far (e.g., majority of the usage analytics studies and the ones in Table 1 are MOOC courses in which the video playback is the central vehicle for content delivery). The interpretation and the application of our findings and discussions on the behavior of the visualization and its effects for other courses, therefore, will need to take this into account.

6.3. Future studies

This study was the first deployment using the strategy that weighed individual interactions differently. One reason for pushing for full deployments integrated into an existing course was to identify which aspects of the system could or should be the main focus of future studies within the design, implementation and deployment, and also how the system and the way the course is delivered could more tightly support each other.

There is room for improvement by incorporating other more detailed playback contexts into the playback weighting. For example, the faster/slower speed setting could have more elaborate weighting variations depending on the target speed set, informed by other studies on the particular speed affecting the students' overall performance; other interaction factors that could be used include the volume level used while playing and whether it was viewed full-screen.

While in this paper we focused on understanding the overall usefulness of the visualization and its potential value, more detailed analysis specifically delving into different categories of session behaviors or student group behaviors with statistical means will allow us to characterize its usefulness in more specific ways. Similarly, designing a comparative study which has a half of the student user population with the visualization and the other half without it or with a different weighting strategy will help understand the level of usefulness of the visualization and its weighting parameters.

One of the potential issues inherent in our approach is when each video is first uploaded and becomes available on the system without any prior playback usage data. The "cold-start" issue is typical of many usage-based machine learning algorithms such as recommender systems (Kang et al., 2019; Natarajan et al., 2020) where the effectiveness of the approach relies on its usage over time. In our system, when a video first becomes available, the visualization does not have usage data from which to calculate and populate the timeline. As soon as a student starts using a video, its usage will be captured and reflected as a low, yellow area at the bottom of the timeline. A manual indication of the important parts by the lecturer at the time of uploading each video may be a simple solution that could also serve as a way to lead a more desirable evolution of the contour. However, given one of the strengths of our approach is fully-automated shaping of the visualization without extra efforts by a human, other more automated ways to draw the data will be an important topic of further study. For example, some studies used supervised machine learning techniques applied to lecture video images to determine the important parts in videos (Brooks et al., 2009) or applied to the usage data (selecting a video, seeking within the video, as well as the heartbeat rates of the students) to develop a computational model of the usage patterns of students to predict a more desirable pattern (Brooks et al., 2014). Taking a bottom-up indexing approach (i.e., important parts emerge by looking at the video/usage data) rather than top-down (i.e., system first defines and codes in what the important events are), bringing some of these computational approaches into our system may help tackle this issue in a way aligned to what we eventually envision. In our second deployment, the same videos used in the first deployment were uploaded afresh without any use of the data captured from the first deployment. One possible approach would be to plug in the usage data from a previous running of the course [if it has run in previous year(s)] at the start of the course, to be slowly overridden by the actual usage once this was sufficient.

More refined versions with a tailored playback weighting strategy could well be applied outside the educational domain. For example, such an indication could be useful cues in general video platforms such as YouTube and Vimeo in which the user communities already drive the access and usage.

7. Conclusion

We introduced visualization guidance to students playing educational videos based on calculating a set of playback weightings for different video interactions from students occurring in different playback contexts. The result shows that for 73.1% of cases, those parts guided by the visualization and thus actually watched by students, were the parts that the instructor judged as the most important parts of the videos. By placing the results of video playback usage not to the post-usage summative stage but to the middle of usage where it will help the users use the system, we broadened the boundaries of video usage analytics, starting with an educational setting where more explicit within-video guidance is very much needed.

Data availability statement

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation. The data used for this study is available on Figshare (DOI: 10.6084/m9.figshare.20288763.v2).

Ethics statement

The studies involving human participants were reviewed and approved by Data Protection Office, Dublin City University. Written informed consent for participation was not required for this study in accordance with the national legislation and the institutional requirements.

Author contributions

HL: conception of the research, UI design, end user issues, discussions, and write up. AS: conception of the research, shaping of all aspects of the study, data analytics, discussions, leading the study, and write up. MS: conception of the research, system implementation, data analytics, discussions, data collection, and write up. ML: conception of the research, system backend, data analytics, system architecture issues, discussions, and write up. All authors contributed to the article and approved the submitted version.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

Atapattu, T., and Falkner, K. (2017). "Discourse analysis to improve the effective engagement of MOOC videos," in *Proceedings of the Seventh International Learning Analytics & Knowledge Conference, LAK '17* (New York, NY: Association for Computing Machinery), 580–581. doi: 10.1145/3027385.3029470

Bentley, F., and Murray, J. (2016). "Understanding video rewatching experiences," in *Proceedings of the ACM International Conference on Interactive Experiences for TV and Online Video, TVX '16* (New York, NY: Association for Computing Machinery), 69–75. doi: 10.1145/2932206.29 32213

Brooks, C., Amundson, K., and Greer, J. (2009). "Detecting significant events in lecture video using supervised machine learning," in *Proceedings of the 14th International Conference on Artificial Intelligence in Education (AIED 2009), Vol. 200* (Brighton), 483–490.

Brooks, C., Erickson, G., Greer, J., and Gutwin, C. (2014). Modelling and quantifying the behaviours of students in lecture capture environments. *Comput. Educ.* 75, 282–292. doi: 10.1016/j.compedu.2014.03.002

Carlier, A., Charvillat, V., and Ooi, W. T. (2015). "A video timeline with bookmarks and prefetch state for faster video browsing," in *Proceedings of the 23rd ACM International Conference on Multimedia, MM '15* (New York, NY: Association for Computing Machinery), 967–970. doi: 10.1145/2733373. 2806376

Chen, T., Peng, L., Yin, X., Rong, J., Yang, J., and Cong, G. (2020). Analysis of user satisfaction with online education platforms in China during the COVID-19 pandemic. *Healthcare* 8,200. doi: 10.3390/healthcare8030200

Choi, S.-M., Lee, H., Han, Y.-S., Man, K. L., and Chong, W. K. (2015). A recommendation model using the bandwagon effect for e-marketing purposes in IOT. *Int. J. Distrib. Sensor Netw.* 11,475163. doi: 10.1155/2015/ 475163

Coman, C., Țîru, L. G., Meseșan-Schmitz, L., Stanciu, C., and Bularca, M. C. (2020). Online teaching and learning in higher education during the coronavirus pandemic: Students' perspective. *Sustainability* 12,10367. doi: 10.3390/su122410367

Cooke, E., Ferguson, P., Gaughan, G., Gurrin, C., Jones, G., Le Borgne, H., et al. (2004). "TRECVid 2004 experiments in Dublin City university," in *TRECVID 2004 - Text REtrieval Conference TRECVID Workshop* (Gaithersburg, MD).

Dawson, S., Joksimovic, S., Poquet, O., and Siemens, G. (2019). "Increasing the impact of learning analytics," in *Proceedings of the 9th International Conference on Learning Analytics & Knowledge, LAK19* (New York, NY: Association for Computing Machinery), 446–455. doi: 10.1145/3303772.33 03784

Dazo, S. L., Stepanek, N. R., Fulkerson, R., and Dorn, B. (2016). "An empirical analysis of video viewing behaviors in flipped CS1 courses," in *Proceedings of the 2016 ACM Conference on Innovation and Technology in Computer Science Education, ITiCSE '16* (New York, NY: Association for Computing Machinery), 106–111. doi: 10.1145/2899415.28 99468

Fong, M., Dodson, S., Harandi, N. M., Seo, K., Yoon, D., Roll, I., and Fels, S. (2019). "Instructors desire student activity, literacy, and video quality analytics to improve video-based blended courses," in *Proceedings of the Sixth (2019) ACM Conference on Learning Scale, L@S '19* (New York, NY: Association for Computing Machinery), 106–111. doi: 10.1145/3330430.33 33618

Gajos, K. Z., Kim, J., Li, S.-W., Cai, C. J., and Miller, R. C. (2014). "Leveraging video interaction data and content analysis to improve video learning," in *Proceedings of the CHI 2014 Learning Innovation at Scale Workshop* (New York, NY: Association for Computing Machinery).

Gaughan, G., Smeaton, A. F., Gurrin, C., Lee, H., and McDonald, K. (2003). "Design, implementation and testing of an interactive video retrieval system," in *Proceedings of the 5th ACM SIGMM International Workshop on Multimedia Information Retrieval, MIR '03* (New York, NY: Association for Computing Machinery), 23–30. doi: 10.1145/973264. 973269

Giannakos, M. N., Chorianopoulos, K., and Chrisochoides, N. (2015). Making sense of video analytics: lessons learned from clickstream interactions, attitudes, and learning outcome in a video-assisted course. *Int. Rev. Res. Open Distrib. Learn*. 16, 260–283. doi: 10.19173/irrodl.v16i1. 1976

Guo, P. J., Kim, J., and Rubin, R. (2014). "How video production affects student engagement: an empirical study of MOOC videos," in *Proceedings of the First ACM Conference on Learning Scale Conference, L@S '14* (New York, NY: Association for Computing Machinery), 41–50. doi: 10.1145/2556325.25 66239

Ihantola, P., Vihavainen, A., Ahadi, A., Butler, M., Börstler, J., Edwards, S. H., et al. (2015). "Educational data mining and learning analytics in programming: literature review and case studies," in *Proceedings of the 2015 ITiCSE on Working Group Reports, ITICSE-WGR '15* (New York, NY: Association for Computing Machinery), 41–63. doi: 10.1145/2858796.285 8798

Joordens, S., Le, A., Grinnell, R., and Chrysostomou, S. (2009). Eating your lectures and having them too: is online lecture availability especially helpful in "skills-based" courses? *Electron. J. e-Learn*. 7, 281–288. Available online at: https:// academic-publishing.org/index.php/ejel/article/view/1577

Jovanović, J., Gašević, D., Dawson, S., Pardo, A., and Mirriahi, N. (2017). Learning analytics to unveil learning strategies in a flipped classroom. *Intern. High. Educ*. 33, 74–85. doi: 10.1016/j.iheduc.2017.0 2.001

Kang, S., Hwang, J., Lee, D., and Yu, H. (2019). "Semi-supervised learning for cross-domain recommendation to cold-start users," in *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM '19* (New York, NY: Association for Computing Machinery), 1563–1572. doi: 10.1145/3357384.335 7914

Kim, J., Guo, P. J., Seaton, D. T., Mitros, P., Gajos, K. Z., and Miller, R. C. (2014a). "Understanding in-video dropouts and interaction peaks inonline lecture videos," in *Proceedings of the First ACM Conference on Learning @ Scale Conference, L@S '14* (New York, NY: Association for Computing Machinery), 31–40. doi: 10.1145/2556325.256 6237

Kim, J., Nguyen, P. T., Weir, S., Guo, P. J., Miller, R. C., and Gajos, K. Z. (2014b). "Crowdsourcing step-by-step information extraction to enhance existing how-to videos," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '14* (New York, NY: Association for Computing Machinery), 4017–4026. doi: 10.1145/2556288.255 6986

Lang, D., Chen, G., Mirzaei, K., and Paepcke, A. (2020). "Is faster better? A study of video playback speed," in *Proceedings of the Tenth International Conference on Learning Analytics & Knowledge, LAK '20* (New York, NY: Association for Computing Machinery), 260–269. doi: 10.1145/3375462.337 5466

Le, A., Joordens, S., Chrysostomou, S., and Grinnell, R. (2010). Online lecture accessibility and its influence on performance in skills-based courses. *Comput. Educ*. 55, 313–319. doi: 10.1016/j.compedu.2010.0 1.017

Lee, H., Balasubramanian, G., and King Wang, P. (2017). "Edubang: envisioning the next generation video sharing," in *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems, CHI EA '17* (New York, NY: Association for Computing Machinery), 454. doi: 10.1145/3027063.304 9775

Lee, H., Liu, M., Riaz, H., Rajasekaren, N., Scriney, M., and Smeaton, A. F. (2021). "Attention based video summaries of live online Zoom classes," in *Proceedings of the 35th AAAI Conference on Artificial Intelligence, Workshop on Imagining Post-COVID Education With Artificial Intelligence*.

Lee, H., Smeaton, A. F., O'Connor, N. E., and Smyth, B. (2006). User evaluation of Físchlár-News: an automatic broadcast news delivery system. *ACM Trans. Inform. Syst*. 24, 145–189. doi: 10.1145/1148020.114 8021

Lehane, B., O'Connor, N. E., Smeaton, A. F., and Lee, H. (2006). "A system for event-based film browsing," in *Technologies for Interactive Digital Storytelling and Entertainment*, eds S. Göbel, R. Malkewitz, and I. Iurgel (Berlin; Heidelberg: Springer Berlin Heidelberg), 334–345. doi: 10.1007/119445 77_33

Leibetseder, A., Kletz, S., and Schoeffmann, K. (2018). "Sketch-based similarity search for collaborative feature maps," in *MultiMedia Modeling*, eds K. Schoeffmann, T. H. Chalidabhongse, C. W. Ngo, S. Aramvith, N. E. O'Connor, Y. S. Ho, M. Gabbouj, and A. Elgammal (Cham: Springer International Publishing), 425–430. doi: 10.1007/978-3-319-7360 0-6_45

Li, N., Kidziński, L., Jermann, P., and Dillenbourg, P. (2015a). "How do in-video interactions reflect perceived video difficulty?," in *Proceedings of the European MOOCs Stakeholder Summit 2015*, 112–121.

Li, N., Kidziński, L., Jermann, P., and Dillenbourg, P. (2015b). "MOOC video interaction patterns: What do they tell us? Design for teaching and learning in a networked world," in *10th European Conference on Technology Enhanced Learning, EC-TEL 2015* (Toledo). doi: 10.1007/978-3-319-2425 8-3_15

Mo, C.-Y., Wang, C., Dai, J., and Jin, P. (2022). Video playback speed influence on learning effect from the perspective of personalized adaptive learning: a study based on cognitive load theory. *Front. Psychol*. 13,839982. doi: 10.3389/fpsyg.2022.839982

Moore, C., Battestilli, L., and Domínguez, I. X. (2021). "Finding video-watching behavior patterns in a flipped CS1 course," in *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education* (New York: Association for Computing Machinery), 768–774. doi: 10.1145/3408877.343 2359

Müller, S., Miller, G., and Fels, S. (2010). "Using temporal video annotation as a navigational aid for video browsing," in *Adjunct Proceedings of the 23nd Annual ACM Symposium on User Interface Software and Technology, UIST '10* (New York, NY: Association for Computing Machinery), 445–446. doi: 10.1145/1866218.186 6263

Natarajan, S., Vairavasundaram, S., Natarajan, S., and Gandomi, A. H. (2020). Resolving data sparsity and cold start problem in collaborative filtering recommender system using linked open data. *Expert Syst. Appl*. 149,113248. doi: 10.1016/j.eswa.2020.113248

Pham, P., and Wang, J. (2015). "Attentivelearner: Improving mobile MOOC learning via implicit heart rate tracking," in *Proceedings of the 17th International Conference on Artificial Intelligence in Education (AIED 2015)* (Madrid). doi: 10.1007/978-3-319-19773-9_37

Popa, D., Repanovici, A., Lupu, D., Norel, M., and Coman, C. (2020). Using mixed methods to understand teaching and learning in COVID 19 times. *Sustainability* 12, 1–20. doi: 10.3390/su122 08726

Robal, T., Zhao, Y., Lofi, C., and Hauff, C. (2018). "IntelliEye: enhancing MOOC learners' video watching experience through real-time attention tracking," in *Proceedings of the 29th on Hypertext and Social Media, HT '18* (New York, NY: Association for Computing Machinery), 106–114. doi: 10.1145/3209542.3209547

Sauter, L., Amiri Parian, M., Gasser, R., Heller, S., Rossetto, L., and Schuldt, H. (2020). "Combining Boolean and multimedia retrieval in vitrivr for large-scale video search," in *MultiMedia Modeling*, eds Y. M. Ro, W. H. Cheng, J. Kim, W. T. Chu, P. Cui, J. W. Choi, M. C. Hu, and W. De Neve (Cham:

Springer International Publishing), 760–765. doi: 10.1007/978-3-030-37734-2_66

Schoeffmann, K. (2019). "Video browser showdown 2012-2019: a review," in *2019 International Conference on Content-Based Multimedia Indexing (CBMI)*, 1–4. doi: 10.1109/CBMI.2019.8877397

Sinha, T., Jermann, P., Li, N., and Dillenbourg, P. (2014). Your click decides your fate: leveraging clickstream patterns in MOOC videos to infer students' information processing and attrition behavior. *arXiv[Preprint].arXiv:1407.7131*. doi: 10.48550/ARXIV.1407.7143

Srivastava, N., Nawaz, S., Newn, J., Lodge, J., Velloso, E., M. Erfani, S., et al. (2021). "Are you with me? Measurement of learners' video-watching attention with eye tracking," in *LAK21: 11th International Learning Analytics and Knowledge Conference, LAK21* (Association for Computing Machinery), 88–98. doi: 10.1145/3448139.3448148

Sundar, S. S., Oeldorf-Hirsch, A., and Xu, Q. (2008). "The bandwagon effect of collaborative filtering technology," in *CHI '08 Extended Abstracts on Human Factors in Computing Systems, CHI EA '08* (New York, NY: Association for Computing Machinery), 3453–3458. doi: 10.1145/1358628.1358873

Wang, K. C., Lai, C.-M., Wang, T., and Wu, S. F. (2015). "Bandwagon effect in facebook discussion groups," in *Proceedings of the ASE BigData & SocialInformatics 2015, ASE BD&SI '15* (New York, NY: Association for Computing Machinery).

Yusuf, N., and Al-Banawi, N. (2013). The impact of changing technology: the case of e-learning. *Contemp. Issues Educ. Res.* 6, 173–180. doi: 10.19030/cier.v6i2.7726