
Distinguishing a ‘hit’ from a ‘view’: Using the access durations of lecture recordings to tell whether learning might have happened

David C. Simcock

James Cook University, Australia
Deviot Institute, Australia

Wei-Hang Chua

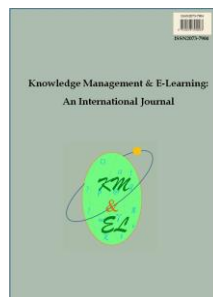
Margreet Hekman

Matthew T. Levin

Massey University, New Zealand

Simon Brown

Deviot Institute, Australia
James Cook University, Australia



Knowledge Management & E-Learning: An International Journal (KM&EL)
ISSN 2073-7904

Recommended citation:

Simcock, D. C., Chua, W. H., Hekman, M., Levin, M. T., & Brown, S. (2019). Distinguishing a ‘hit’ from a ‘view’: Using the access durations of lecture recordings to tell whether learning might have happened. *Knowledge Management & E-Learning*, 11(1), 81–94.

Distinguishing a ‘hit’ from a ‘view’: Using the access durations of lecture recordings to tell whether learning might have happened

David C. Simcock

College of Public Health, Medical and Veterinary Sciences
James Cook University, Australia
Deviot Institute, Australia
E-mail: David.Simcock@jcu.edu.au

Wei-Hang Chua 

Institute of Food Science and Technology
Massey University, New Zealand
E-mail: W.H.Chua@massey.ac.nz

Margreet Hekman 

Institute of Veterinary, Animal and Biomedical Science
Massey University, New Zealand
E-mail: M.Hekman@massey.ac.nz

Matthew T. Levin

Institute of Food Science and Technology
Massey University, New Zealand
E-mail: M.T.Levin@massey.ac.nz

Simon Brown* 

Deviot Institute, Australia
College of Public Health, Medical and Veterinary Sciences
James Cook University, Australia
E-mail: Simon.Brown@deviotinstitute.org

*Corresponding author

Abstract: Audiovisual recordings of lectures are available to many students in all disciplines. The use of lecture recordings has been studied extensively, but it is still not clear how, or how much, they are actually used. Previous analysis of their use has been based on either survey data or computer logs of access. In the latter case, measurements of actual use have usually been based on counts of the number of times recordings have been accessed. This does not distinguish those that happen accidentally (‘hits’), from those that might permit learning (‘views’). This distinction is essential to the meaningful analysis of the log of the actual use of recorded lectures. Using the access logs of undergraduate

science students, we show that the distribution of the durations of the access of recordings of scheduled lectures has two distinct components. The most rapid of these is complete within three minutes and we infer that it reflects the behaviour of students searching among recordings. This inference is based on a comparison of these distributions with those of (i) recordings made automatically during a non-teaching period and (ii) individual users. This is also consistent with the pattern of usage by students searching for a specific recording.

Keywords: Online learning; Recorded lectures; Science education; Weibull distribution

Biographical notes: David Simcock is a physiologist with interests in biochemical parasitology, pathophysiology and learning patterns of students studying science in tertiary education. After completing his PhD at Massey University, he has taught large courses for students in science, health science and veterinary science programmes at Massey University in New Zealand and James Cook University in Australia.

Wei-Hang Chua is a physiologist with research interests in bone cell biology and ion channel function. Since completing his PhD, he has conducted research in the field of bone biology and taught and coordinated courses in physiology for science and veterinary science students at Massey University.

Margreet Hekman is an animal scientist completing a PhD at Massey University in New Zealand. She teaches physiology and animal science at Massey University.

Matthew Levin is an IT consultant and service support manager at Massey University. He is interested in usage patterns in online media resources in education.

Simon Brown is a biochemist with interests in bioenergetics, mathematical analysis and science education. After completing a PhD at the Australian National University, he taught and carried out research at universities in the United Kingdom, New Zealand and Australia. He is now at the Deviot Institute and James Cook University.

1. Introduction

Audiovisual recordings of ordinary lectures are increasingly common in tertiary education, especially for large classes (Owston, Lupshenyuk, & Wideman, 2011). The recordings are usually available only through systems administered by the institution concerned using software such as Moodle or Blackboard. These recordings are intended to help students, so how they are used and whether they actually are helpful is of interest to their teachers (and, perhaps, to the administrators of the institution). So, it is of some concern that only 51% of the mathematics students surveyed by Yoon, Oates, and Sneddon (2014) intended to make use of lecture recordings and only 52% of respondents in a recent survey of undergraduate science students claimed to have actually done so (Simcock et al., 2017). While this is consistent with the preference of many students for a live teacher in the room rather than a recorded lecture (Simcock et al., 2017), these estimates are based on students' own reports and it is important to assess their reliability using computer records of their actual usage.

The first step in analysing the record of actual use of lecture recordings in the latter case (Simcock et al., 2017) is to address the fundamental question: what constitutes a useful ‘view’ of a recording? Here, we distinguish between a brief access of a lecture recording (a ‘hit’), which might occur when a recording is selected in error or when a user rapidly realises that a recording is not what was being sought, and a longer ‘view’, in which there is a reasonable chance that some learning could occur. It is conceivable that learning might occur in some ‘hits’ and that it might not occur in some ‘views’, but without other information about student engagement with the recording, such as might be available from research tools facilitating the analysis of ‘clickstream’ data (Brooks, Greer, & Gutwin, 2014) or physiological data (Chen & Wu, 2015), it is not possible to distinguish these.

If the value of lecture recordings is to be properly assessed, it is important to be able to distinguish a potentially useful ‘view’ of a lecture recording from a mere ‘hit’ using simple, readily available data. Many reports are based on surveys in which students are asked to estimate their own usage (Azab et al., 2016; Dommeyer, 2017). Where this is not the case, previous work on the use of lecture recordings has been based on three different perspectives of the relationship between a ‘hit’ and a ‘view’. First, several reports are based on the assumption that whenever a recording is ‘accessed’ it is used for learning, no matter the duration (Danielson et al., 2014; Dickson et al., 2012; Leadbeater et al., 2013; Mark & Vrijmoed, 2016; Owston, Lupshenyuk, & Wideman, 2011; Ozan & Ozarslan, 2016). In fact, information concerning the duration of access is given in only two of these reports (Mark & Vrijmoed, 2016; Ozan & Ozarslan, 2016). A slightly different position is adopted by Johnston, Massa, and Burne (2013) who suggest that most, rather than every, access constituted a ‘view’. However, they did not state how they came to this conclusion, other than stating that their analysis was ‘qualitative’, nor did they explain how one might distinguish a ‘hit’ from a ‘view’. Second, in at least one report in which ‘hits’ are specified, it is explicitly acknowledged that there is no evidence that each ‘hit’ constitutes a ‘view’ (Williams, Pfeifer, & Waller, 2013). Third, Marchand, Pearson, and Albon (2014) reported the number of accesses, total viewing times and the range of access durations, but they did not report the distribution of the latter or make any attempt to analyse these data further. In order to reduce the “novelty-effect bias associated with having a new tool in the learning environment”, Brooks, Erickson, Greer, and Gutwin (2014) distinguished between those students who “... watched at least five minutes of video lecture content in a calendar week” and those who had not, and the latter were deemed not to have watched any content. While the reason given for this approach is quite different, it goes some way towards distinguishing between ‘hits’ and ‘views’. These observations prompt the question that motivates the work we describe, which is whether there is an objective means of distinguishing a useful ‘view’ of a lecture recording from a ‘hit’ using simple, readily available data.

This question is not unique to lecture recordings. It is closely related to the more general problem of distinguishing an ‘event’ from a mere ‘attempt’. For example, the abandonment of a view of a lecture recording is akin to a caller ending a telephone call (either before or after the call is answered) (Gans, Koole, & Mandelbaum, 2003; Jiang et al., 2013), a web surfer moving on to the next web page (Liu, White, & Dumais, 2010) or the eyes moving on to the next word when reading (Feng, 2009). To make this analogy explicit, consider that any access of a lecture recording involves at least three essential user-initiated acts: (i) selection of a recording, (ii) initiation of access and (iii) termination of access. The time between the initiation and termination is the duration. While information other than these five variables (the user, identity of the target recording, start and end times, and duration) might be available in some circumstances (Brooks, Greer, &

Gutwin, 2014), these represent the essence of the process. The same five variables (the caller, the identity of the target number, start and end times, and duration) characterise a telephone call in which a particular caller (i) selects a number to call, (ii) initiates a call and (iii) terminates the call. In neither case is any information available about events that occur between initiation and termination. A telephone call is analogous to accessing a lecture recording access in one other significant respect: the identity of the person actually initiating the event is usually, but need not always be, the ‘owner’ of the telephone number from which a call is placed or the user account employed to access a lecture recording. Other information might be available, such as the provider of the telephone or internet service, but such data are not an essential characteristic of the event in question because there is no reason to expect that it would be changed if it happened to be different. Nevertheless, ‘extraneous’ data of this type extend the analogy between placing a telephone call and accessing a lecture recording.

In both placing a telephone call and accessing a lecture recording the attempt can have been useful only if the time before termination (the duration) is sufficient, but is unlikely to have been if the duration is very brief. The opportunity for information transfer increases with the duration of the event, if only because more can be said in 10 minutes than in 10 seconds, whether in a telephone call or a recorded lecture. Of course, a very small proportion of telephone calls could be terminated after 10 seconds without any loss of information, but this is rarely the case for lectures. It follows from this that the durations provide some insight into the likelihood of information transfer. For example, even if a telephone call goes unanswered, the caller can conclude that there will not be a response if it is allowed to ring for long enough, but no reliable inference about the likelihood of a response can be made if the call is terminated so rapidly that it could not have been answered. In this case, the duration of the call is usually treated as though it has the Weibull distribution, for which there is some quasi-theoretical justification. For example, Palm (1953) related the duration of a telephone call to the inconvenience (I) of the caller and modelled the derivative of I (which he called the irritation) as a power function ($dI = ct^k dt$, where c is a constant, t is the elapsed time and k represents the strength of the relationship) so that the irritation (and the inconvenience) increases with the duration. If the irritation is proportional to the hazard rate of abandonment, then the duration of telephone calls has the Weibull distribution (Weibull, 1951). The duration data also encode information about the behaviour of different users, as Palm (1953) appreciated, and we consider some aspects of this.

Unlike telephone calls, the record of viewing of recorded lectures has the advantage that both the pattern of use by individuals and the pattern of use of individual recordings can be analysed. For example, one student might know precisely which lecture to access and simply watch it in full or in part, but another might have only a vague idea which lecture to view. The latter individual is likely to exhibit behaviours that reflect searching or frustration, as well as viewing. An examination of these sorts of behaviours can be used to interpret the significance of the data (Feild, Allan, & Jones, 2010; Liu, White, & Dumais, 2010; Wang, Lin, & Chen, 2010). Nevertheless, in analysing any online resource usage it is necessary to distinguish a ‘view’ from a mere ‘hit’. Here, we describe and analyse users’ search behaviours and consider those features that distinguish a ‘view’ of a recorded lecture from a ‘hit’.

2. Methods

2.1. Data

As described previously (Simcock et al., 2017), students enrolled in a one semester course entitled ‘Essentials of mammalian biology’ were asked about their use of recorded lectures and their actual usage was recorded automatically by the university using Mediasite. The survey, the collection of online data and the protocol employed were approved by the Massey University Human Ethics Committee (B – southern North Island) and the Head of the College of Science, Massey University. The lecture recordings were provided to the students as a resource for them to use as they wished and without any particular suggestion as to how they might be used. Some data about how students claimed to use these lecture recordings have been reported previously (Simcock et al., 2017).

The actual usage records of the 145 users (of a total of 267 students enrolled) who participated in the survey (Simcock et al., 2017) and consented to the use of their data were extracted from the Mediasite log. Of these, 96 users accessed video recordings of lectures a total of 1866 times and the recorded duration of these hits ranged from 1 s (the resolution of the reported measurement) to 23.52 h (Table 1). The number of hits ranged from 29 to 86 for each lecture recording. Two other features of the data should be noted. First, in the middle of each semester teaching stops for about a week. No lectures were delivered during the mid-semester break, but the system automatically made recordings at the scheduled lecture times (we refer to these as the ‘break’ lectures). These recordings were loaded onto the server and were accessed by some users (Table 1). Second, a student played a practical joke (a ‘prank’) during the last lecture (which we refer to as the ‘prank’ lecture). The prank was very brief, but it excited some interest among the students and prompted a larger number of users (86) to access the recording than was the case for previous lectures (Table 1).

Table 1
Properties of the distributions of the views of recorded lectures

	Recordings			
	All	All except ‘prank’ and ‘break’ lectures	‘break’	‘prank’
number of				
views	1866	1631	149	86
recordings	40	35	4	1
users	96	86	39	60
mean (h)	1.28	1.35	0.27	1.54
SD (h)	2.21	2.28	0.80	2.21
median (h)	0.59	0.77	0.01	0.74
IQR (h)	0.01-2.39	0.01-23.53	0.00-0.02	0.09-2.58
range (h)	0.00-23.53	0.00-23.53	0.00-2.75	0.00-17.06

2.2. Analysis

Each time a user started or stopped accessing a lecture recording the event was recorded in the Mediasite log. The duration of an access (t , in hours), calculated from these times, is analogous to the duration of a telephone call (Palm, 1953) and can be considered to have a Weibull distribution (Weibull, 1951). The Weibull distribution is a very common choice when considering measurements of duration derived from processes that share some of the characteristics identified by Palm (1953) in his treatment of telephone calls (Bučar, Nagode, & Fajdiga, 2004; Feng, 2009; Gans, Koole, & Mandelbaum, 2003; Jiang et al., 2013; Liu, White, & Dumais, 2010; Razali & Al-Wakeel, 2013). We also tested mixtures of the gamma distribution, but this yielded a poorer fit to the data.

The probability density function (PDF) of the Weibull distribution is

$$f(t; k, \lambda) = \frac{k}{\lambda} \left(\frac{t}{\lambda} \right)^{k-1} \exp \left[- \left(\frac{t}{\lambda} \right)^k \right], \quad t > 0, k > 0, \lambda > 0, \quad (1)$$

where k (dimensionless) and λ (in hours) are the ‘shape’ and ‘scale’ parameters, respectively. The magnitude of k changes shape of the distribution:

- i. if $k < 1$, the PDF is large as t approaches zero and tends towards zero as t increases;
- ii. for $k = 1$, the Weibull distribution is identical to the exponential distribution and the PDF approaches λ^{-1} as t approaches 0; and
- iii. for $k > 1$ the PDF is low at small t , rises, passes through a maximum and then decreases as t increases.

The ‘scale’ parameter (λ) determines the value of t at which $f(t) \approx 0.632k/\lambda$ and the mean and variance of a Weibull variable (t in this case) are proportional to λ and λ^2 , respectively. The corresponding cumulative distribution function (CDF) is

$$F(t; k, \lambda) = 1 - \exp \left[- \left(\frac{t}{\lambda} \right)^k \right], \quad (2)$$

where $1-F(t)$ is often called the reliability (R). In essence, the effects of a larger (smaller) λ is to move $F(t)$ to higher (lower) t and to increase (decrease) the steepness of the curve, and a larger (smaller) k tends to make $F(t)$ increase more (less) steeply with increasing t . The hazard (or failure) rate is

$$h(t; k, \lambda) = \lim_{\Delta t \rightarrow 0} \frac{P(t \leq X < t + \Delta t \mid X \geq t)}{\Delta t} = \frac{f(t)}{1 - F(t)}, \quad (3)$$

and that corresponding to (1) and (2) is

$$h(t; k, \lambda) = \frac{k}{\lambda} \left(\frac{t}{\lambda} \right)^{k-1}, \quad (4)$$

which has the dimension of h^{-1} here. If $k < 1$ or $k > 1$ the hazard rate declines or increases, respectively, as t increases, and if $k = 1$ the hazard rate is constant.

In complex systems it is often necessary to combine two or more Weibull components in order to account for the distribution (Woodward & Gunst, 1987). For two components (1) becomes

$$f_2(t) = p_1 f(t; k_1, \lambda_1) + (1 - p_1) f(t; k_2, \lambda_2), \quad (5)$$

where $0 \leq p_1 \leq 1$ is the contribution of $f(t; k_1, \lambda_1)$ and two shape (k_1 and k_2) and two scale (λ_1 and λ_2) parameters are also required, and (2) becomes

$$F_2(t) = 1 - p_1 \exp\left[-\left(\frac{t}{\lambda_1}\right)^{k_1}\right] - (1 - p_1) \exp\left[-\left(\frac{t}{\lambda_2}\right)^{k_2}\right]. \quad (6)$$

Analogous expressions for n components follow directly from these (Bučar, Nagode, & Fajdiga, 2004; Davison & Louzada-Neto, 2000; Panteleeva, Gutiérrez González, Vaquera Huerta & Villaseñor Alva, 2015; Razali & Al-Wakeel, 2013) and the corresponding hazard rate can be calculated from (3). By choosing mixtures of the Weibull distribution we do not intend to imply that this is the only possibility. We merely make the point that it has some pseudo-theoretical justification (Palm, 1953) and that it provides a good fit to the data.

All analyses were carried out in R (Ihaka & Gentleman, 1996) and hazard rates were estimated from the data using the *muhaz* package.

3. Results

3.1. Overall distribution of access durations

The empirical CDF of the access durations (Fig. 1A) indicates that there were at least four components:

- i. a rapid phase ($k_1 = 0.675 \pm 0.009$, $\lambda_1 = 0.0085 \pm 0.0001$ h, $p_1 = 0.397$),
- ii. a phase centred at about 1 h ($k_2 = 8.6 \pm 0.6$, $\lambda_2 = 0.933 \pm 0.006$ h, $p_2 = 0.068$),
- iii. a prolonged phase ($k_3 = 1.01 \pm 0.02$, $\lambda_3 = 1.63 \pm 0.02$ h, $p_3 = 0.342$) and
- iv. an extended phase ($k_4 = 20 \pm 1$, $\lambda_4 = 2.758 \pm 0.004$ h, $p_4 = 0.193$).

It is apparent from Fig. 1A that this mixture of Weibull components diverges from the data for t less than about 0.005 h (or 18 s), but it is unlikely that an access of this duration could be very useful. The hazard function corresponding to the entire dataset is also shown in Fig. 1A. It confirms that there is a rapid decrease in the rate at which accesses are abandoned prior to a short period during which the rate was roughly constant. Subsequently, peaks at 0.9 h and 2.7 h are apparent before a very small increase, corresponding to the small number of very prolonged ‘views’. Phases i, ii and iv are also apparent from the frequency distribution of the access durations (Fig. 1B). The prolonged phase is less obvious, but it is consistent with the baseline frequency apparent between 0.1 h and about 0.6 h (Fig. 1B) that also underlies the distribution for $t < 0.1$ h.

In the first phase $k < 1$, so the rate of abandonment declines, and is essentially completed within 3 min (= 0.05 h) (Fig. 1A). This is likely to be the time required to realise that the wrong recording had been selected and end the process. If this interpretation is correct, the inevitable inference is that about 40% of all events were rapidly (within 3 min) terminated and that a ‘real’ view must last for more than 3 min. The second phase ($k > 1$) may represent those users who watched the entire recording. As we have previously reported (Simcock et al., 2017), a significant number of users (73% of those who used the recordings) reported that this was the way they usually watched

recordings, but this phase accounted for only 6.8% of accesses. In the third phase $k \approx 1$, consistent with an approximately constant rate of abandonment (3) that is usually interpreted as an indication of a random process. This phase accounts for about 35% of accesses. The fourth phase ($k > 1$) was of some concern because it was substantial (19.3%, Fig. 1) and unexpected for views of lecture recordings that lasted less than 1 h because $\lambda \approx 2.8$ h. Such extended views may represent those users who watch the entire recording, but intermittently pause to make notes and view some sections more than once. However, closer inspection of values in this range showed that eleven durations occurred more than 10 times ((t, count) : (2.355000, 10), (2.500000, 12), (2.525000, 13), (2.579167, 31), (2.752500, 71), (2.751667, 94), (2.753333, 13), (3.000833, 15), (3.001667, 31), (3.002500, 32), (3.003333, 12)). It will be apparent that these may well be represented by an even smaller set of approximate durations ((2.4, 10), (2.5, 56), (2.75, 178), (3.0, 90)), which reinforces the speculation that these values might represent users being automatically timed out of viewing sessions, perhaps because of inactivity. That views of as much as 23.53 h (Table 1) were recorded does complicate the interpretation of this, but the discrepancy may reflect differences between the users' internet service providers and the university network.

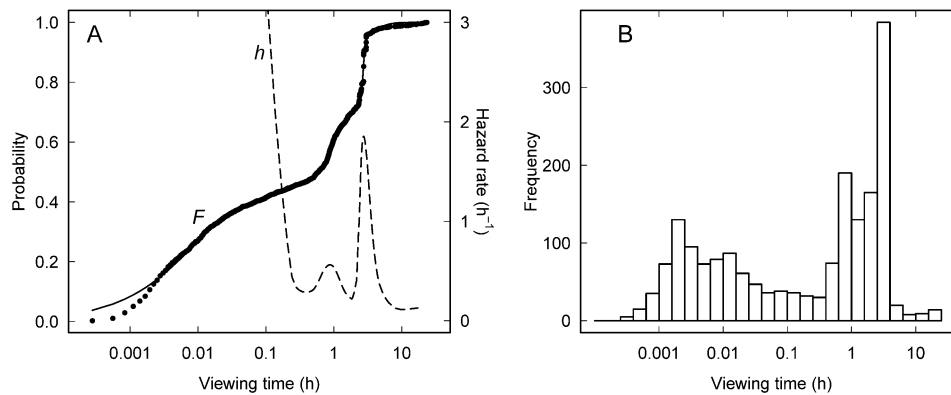


Fig. 1. Cumulative distribution (F) and hazard rate (h) of terminations for (A) and distribution of access durations of (B) the entire dataset (96 users, 1866 accesses). In (A) the solid circles (\bullet) represent the data to which was fitted the four component Weibull mixture (—) described in the text and from which the hazard rate (---) was estimated. A summary of the properties of the distribution is given in Table 1.

3.2. Differences between lectures

The distribution of t for each of the 40 lecture recordings is simpler than that of the more complex empirical CDF for the entire dataset (Fig. 1A). Thirty-six of these recordings were adequately described using two Weibull components (5, 6), and the second component was not necessary for the other four recordings of the ‘break’ lectures (Fig. 2). The latter, such as that shown in Fig. 2, were all fitted to one Weibull component and in each case most of the views were terminated within about 0.05 h (3 min). This is presumably an indication of the time needed by a user to decide that there was nothing to see and terminate the view. This rapid phase is also apparent, to differing extents, in the distributions of the view durations of recordings of scheduled lectures (Fig. 2). The variation in the amplitude of this phase presumably reflects differences between users and in how easy it is to identify that a specific recording is not the one required. The second

phase is often complete by about 1 h, although some last about 2.75 h and a very few last longer than 10 h (Fig. 2).

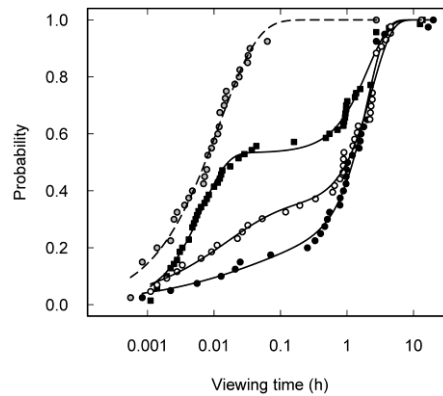


Fig. 2. Cumulative distribution of the accesses of selected lecture recordings. Each symbol indicates a different lecture, one of which was a ‘break’ lecture (●, ---), and each of the curves, except that for the ‘break’ lecture, is a fit of two Weibull components (5) to the corresponding data. Only one Weibull component is fitted to the ‘break’ lecture data.

The hazard rates obtained from the ‘blank’ and ‘prank’ lectures were reminiscent of the ‘bathtub’ failure profile (Klutke, Kiessler, & Wortman, 2003; Wondmagegnehu, Navarro, & Hernandez, 2005), quite different from that in Fig. 1A (data not shown). The ‘blank’ lectures were dominated by the initial phase during which the hazard function declined rapidly, although the small number of 2.75–3 h sessions did generate a slight increase in the hazard rate. After a very brief period during which the hazard rate increased, the ‘prank’ lecture was similar to the ‘blank’ lectures.

The accesses of the ‘break’ and ‘prank’ lecture recordings make a relatively minor contribution to the overall distribution of t (compare Fig. 1B with Fig. 3A). However, the distribution of the views of the 35 standard scheduled lecture recordings is essentially bimodal (Fig. 3A). Most of the accesses of the four ‘break’ lecture recordings were over within 0.1 h (Fig. 3B), consistent with the example shown in Fig. 2, which corresponds to the first peak in the distribution of views of standard scheduled lectures (Fig. 3A). The most frequently accessed recording was that of the ‘prank’ lecture. The distribution of t is compressed in this case because there are fewer of the shortest accesses (the lower quartile was 0.09 h rather than 0.01 h for the standard lecture recordings, but the upper limit was similar to the standard lecture recordings (Table 1)). This is perhaps consistent with users being willing to spend more time searching for the prank than they might for a particular part of a lecture.

3.3. User behaviour

The distribution of t for each user is also simpler than that of the more complex empirical CDF for the entire dataset (Fig. 1A) in that each could be adequately described using two Weibull components (Fig. 4). Understandably, the range of variation among users is considerable. For example, some users had very few short accesses, others have many and then a relatively evenly distributed range of t and there is a great deal of variation

between (Fig. 4). Despite this, the general pattern of a rapid phase and a slower phase centred on 1-2.75 h is apparent.

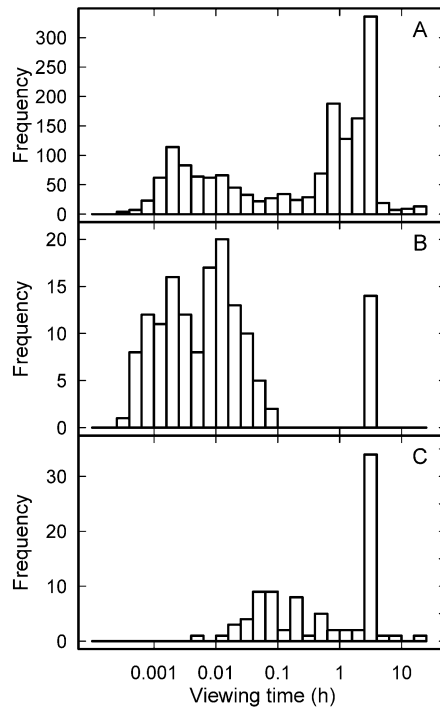


Fig. 3. Distribution of the accesses of (A) all the recordings of scheduled lectures except the ‘prank’ lecture, (B) the ‘break’ lectures and (C) the ‘prank’ lecture. A summary of the properties of each distribution is given in Table 1.

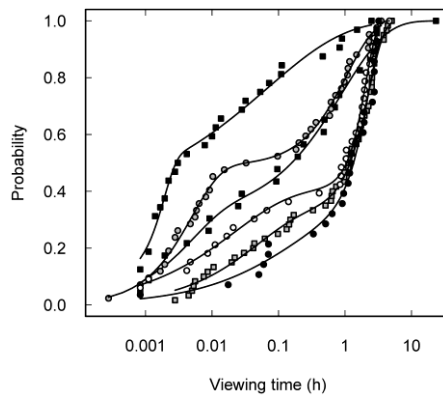


Fig. 4. Cumulative distribution of the duration of accesses of recorded lectures of selected users. Each symbol indicates a different user and the curves are fits of two Weibull components (5) to the data.

Further insight into user behaviour can be obtained from an examination of the details of the record for a single user. For example, searching is very apparent from a repeated sequence of short accesses of less than 0.05 h before a prolonged access of a recording (Fig. 5A). A similar pattern is also apparent from Fig. 5B and, to a lesser extent,

Fig. 5C, which is consistent with the association of the rapid phase of the distribution of t with searching behaviour. The record shown in Fig. 5B also illustrates the fact that a single user can run two simultaneous accesses of a single lecture and the searching prior to the initiation of the second access prompts the speculation that this might have been intentional. A more restricted search pattern and another example of concurrent accesses of the same lecture are apparent in Fig. 5C, which also shows an example of simultaneous accesses of different lectures. The culmination of this period is a prolonged access (of more than 12 h) of lecture 18 that ran overnight, perhaps consistent with the user forgetting to terminate the session.

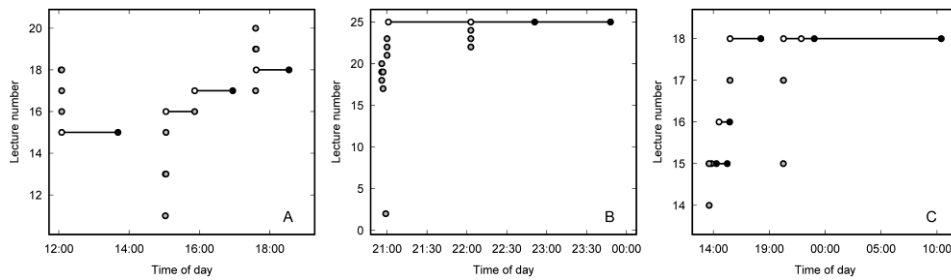


Fig. 5. Portions of the records of individual users for one afternoon (A), one evening (B) and almost one day (C). The initiation and termination of a session are marked by an open circle (○) and solid circle (●, ●), respectively, and in the latter case the grey circle (●) indicates that the view lasted less than 0.05 h. Where only the solid circle is visible it is superimposed on the open circle. Note that lectures 19-22 are the ‘blank’ lectures recorded automatically during the semester break. In (B) the two simultaneous accesses of lecture 25 lasted 2.79 h and 0.79 h, respectively. In (C) the two simultaneous accesses of lecture 15 lasted 1.59 h and 0.46 h and the three simultaneous accesses of lecture 18 lasted 2.75 h, 2.75 h and 12.52 h.

3.4. Reliability of user reports of usage

Of the 145 survey respondents for whom it is possible to identify usage of recorded lectures, 96 (66%) accessed at least one lecture recording (Table 1), which is rather more than the 52% that claimed to have used the lecture recordings (Simcock et al., 2017). However, 14 (15%) of these accessed only one lecture and, of these, 8 users were logged on for less than 20 min (and 4 of these accessed a recording for less than 5 min). From this we infer that these 14 users did not make significant use of the recordings and it follows that it is likely to be reasonable that only 52% of users claimed to have made use of the recordings (Simcock et al., 2017). Irrespective of this, it is remarkable that 34% of respondents did not access even one of the recorded lectures.

4. Discussion

The distribution of the duration of accesses of recorded lectures has two main phases. The first is a rapid phase that is essentially complete within 3 min and the second tended to end at about 1 h or about 2.75 h (Fig. 2 and Fig. 4). Combining all of these data yields a more complex distribution with at least two other components (Fig. 1A). Nevertheless,

the simplest explanation is that the rapid phase corresponds to ‘hits’ that arise from searching behaviour, from which we infer that real ‘views’ lasted more than 3 min.

The choice of the Weibull distribution to describe the data is based on its use with analogous data (Bučar, Nagode, & Fajdiga, 2004; Feng, 2009; Gans, Koole, & Mandelbaum, 2003; Jiang et al., 2013; Liu, White, & Dumais, 2010; Palm, 1953; Razali & Al-Wakeel, 2013) and it provides a convenient framework for interpreting the distributions of the access durations. The best justification for this choice is that the Weibull distribution fits the data well (Fig. 1A, Fig. 2, and Fig. 4), but had we made some other choice (which would have been a poorer fit to the data) it would not have changed the data. For example, the access duration distributions in Fig. 2 (except for that of the ‘break’ lecture) and Fig. 4 have two distinct components in various proportions, which does not depend on the distribution chosen to analyse the data.

The use of lecture recordings has been studied extensively (O’Callaghan et al., 2017) and many reports of their usage are based on the analysis of logs of student access. However, many of these are based on the assumption that whenever a lecture recording is ‘accessed’ it is used for learning, no matter the duration (Danielson et al., 2014; Dickson et al., 2012; Leadbeater et al., 2013; Mark & Vrijmoed, 2016; Owston, Lupshenyuk, & Wideman, 2011; Ozan & Ozarslan, 2016). Even when it is acknowledged that this is not the case (Johnston et al., 2013; Marchand, Pearson, & Albon, 2014; Williams, Pfeifer, & Waller, 2013) no attempt is made to account for the contribution of those that are too short to be a significant opportunity for learning (what we call ‘hits’). The data we report here indicate that these account for about 40% of the total (Fig. 1A). Unsurprisingly, the proportion of ‘hits’ differs among students (Fig. 4) and lectures (Fig. 2), but almost all of the instances of a student accessing one of the ‘break’ lecture recordings (which simply showed an empty lecture theatre) were ‘hits’ lasting less than about 3 min (Fig. 2). If our data are representative of other studies, the 40% ‘hits’ might well obscure, or distort estimates of the magnitude of the effect of lecture recordings on student performance, for example.


We have estimated that a ‘hit’ lasts no longer than about 3 min for this cohort of students. We acknowledge that at least two limitations are associated with this. First, it is possible that this limit might not be similar in other contexts, so it would be interesting to repeat this analysis for other cohorts of students. Second, without some other measure of the engagement of a student with a lecture recording, such as ‘clickstream’ or physiological data (Brooks, Greer, & Gutwin, 2014; Chen & Wu, 2015), there is no way to confirm that our interpretation of ‘hits’ and ‘views’ is correct. It is possible that some students can learn something from a ‘hit’, but, arguably, it is more likely that some students do not learn in some ‘views’. This issue also warrants further consideration.

If the distinction between ‘hits’ and ‘views’ is a reasonable interpretation of the access duration distribution, then about 40% of logged accesses were simply transient ‘hits’ forming part of a search. This means that the usefulness of recorded lectures cannot simply be estimated from the number of times they have been accessed, and combined with the observation that only about 50% of users accessed more than one recording, prompts concern about the real value of recorded lectures.

ORCID

Wei-Hang Chua  <https://orcid.org/0000-0001-5330-8277>

Margreet Hekman  <https://orcid.org/0000-0002-5035-1728>

Simon Brown  <https://orcid.org/0000-0002-0845-4950>

References

- Azab, E., Saksena, Y., Alghanem, T., Midle, J. B., Molgaard, K., Albright, S., & Karimbux, N. (2016). Relationship among dental students' class lecture attendance, use of online resources, and performance. *Journal of Dental Education*, 80(4), 452–458.
- Brooks, C., Erickson, G., Greer, J., & Gutwin, C. (2014). Modelling and quantifying the behaviours of students in lecture capture environments. *Computers & Education*, 75, 282–292.
- Brooks, C., Greer, J., & Gutwin, C. (2014). The data-assisted approach to building intelligent technology-enhanced learning environments. In J. A. Larusson & B. White (Eds.), *Learning Analytics: From Research to Practice* (pp. 123–156). New York, NY: Springer.
- Bučar, T., Nagode, M., & Fajdiga, M. (2004). Reliability approximation using finite Weibull mixture distributions. *Reliability Engineering & System Safety*, 84(3), 241–251.
- Chen, C.-M., & Wu, C.-H. (2015). Effects of different video lecture types on sustained attention, emotion, cognitive load, and learning performance. *Computers & Education*, 80, 108–121.
- Danielson, J., Preast, V., Bender, H., & Hassall, L. (2014). Is the effectiveness of lecture capture related to teaching approach or content type? *Computers & Education*, 72, 121–131.
- Davison, A. C., & Louzada-Neto, F. (2000). Inference for the Poly-Weibull model. *The Statistician*, 49(2), 189–196.
- Dickson, P. E., Warshow, D. I., Goebel, A. C., Roache, C. C., & Adrion, W. R. (2012). Student reactions to classroom lecture capture. In *Proceedings of the 17th ACM Annual Conference on Innovation and Technology in Computer Science Education* (pp. 144–149). Haifa.
- Dommeyer, C. J. (2017). Lecture capturing: Its effects on students' absenteeism, performance, and impressions in a traditional marketing research course. *Journal of Education for Business*, 92(8), 388–395.
- Feild, H. A., Allan, J., & Jones, R. (2010). Predicting searcher frustration. In *Proceedings of the 33rd annual international ACM SIGIR conference on research and development in information retrieval* (pp. 34–41). ACM.
- Feng, G. (2009). Time course and hazard function: a distributional analysis of fixation of duration in reading. *Journal of Eye Movement Research*, 3(2): 2. Retrieved from <https://bop.unibe.ch/index.php/JEMR/article/view/2292>
- Gans, N., Koole, G., & Mandelbaum, A. (2003). Telephone call centers: Tutorial, review, and research prospects. *Manufacturing & Service Operations Management*, 5(2), 79–141.
- Ihaka, R., & Gentleman, R. (1996). R: A language for data analysis and graphics. *Journal of Computational and Graphical Statistics*, 5(3), 299–314.
- Jiang, Z.-Q., Xie, W.-J., Li, M.-X., Podobnik, B., Zhou, W.-X., & Stanley, H. E. (2013). Calling patterns in human communication dynamics. *PNAS*, 110(5), 1600–1605.
- Johnston, A. N. B., Massa, H., & Burne, T. H. J. (2013). Digital lecture recording: A cautionary tale. *Nurse Education in Practice*, 13(1), 40–47.
- Klutke, G.-A., Kiessler, P. C., & Wortman, M. A. (2003). A critical look at the bathtub curve. *IEEE Transactions on Reliability*, 52(1), 125–129.

- Leadbeater, W., Shuttleworth, T., Couperthwaite, J., & Nightingale, K. P. (2013). Evaluating the use and impact of lecture recording in undergraduates: Evidence for distinct approaches by different groups of students. *Computers & Education*, *61*, 185–192.
- Liu, C., White, R. W., & Dumais, S. (2010). Understanding web browsing behaviors through Weibull analysis of dwell times. In *Proceedings of the 33rd annual international ACM SIGIR conference on research and development in information retrieval* (pp. 379–386). ACM.
- Marchand, J.-P., Pearson, M. L., & Albon, S. P. (2014). Student and faculty member perspectives on lecture capture in pharmacy education. *American Journal of Pharmaceutical Education*, *78*(4): 74.
- Mark, K. P., & Vrijmoed, L. L. P. (2016). Does lecture capturing improve learning? A data driven exploratory study on the effectiveness of lecture capture on learning in a foundation IT course. In *Proceedings of the 2016 IEEE international conference on teaching, assessment, and learning for engineering (TALE)*. IEEE.
- O’Callaghan, F. V., Neumann, D. L., Jones, L., & Creed, P. A. (2017). The use of lecture recordings in higher education: a review of institutional, student, and lecturer issues. *Education and Information Technologies*, *22*(1), 399–415.
- Owston, R., Lupshenyuk, D., & Wideman, H. (2011). Lecture capture in large undergraduate classes: Student perceptions and academic performance. *The Internet and Higher Education*, *14*(4), 262–268.
- Ozan, O., & Ozarlan, Y. (2016). Video lecture watching behaviors of learners in online courses. *Educational Media International*, *53*(1), 27–41.
- Palm, C. (1953). Methods of judging the annoyance caused by congestion. *Tele*, *2*, 1–20.
- Panteleeva, O. V., Gutiérrez González, E., Vaquera Huerta, H., & Villaseñor Alva, J. A. (2015). Identifiability and comparison of estimation methods on Weibull mixture models. *Communications in Statistics: Simulation and Computation*, *44*(7), 1879–1900.
- Razali, A. M., & Al-Wakeel, A. A. (2013). Mixture Weibull distributions for fitting failure times data. *Applied Mathematics and Computation*, *219*(24), 11358–11364.
- Simcock, D. C., Chua, W. H., Hekman, M., Levin, M. T., & Brown, S. (2017). A survey of first year biology student opinions regarding live lectures and recorded lectures as learning tools. *Advances in Physiology Education*, *41*(1), 69–76.
- Wang, C.-J., Lin, K. H.-Y., & Chen, H.-H. (2010). Intent boundary detection in search query logs. In *Proceedings of the 33rd annual international ACM SIGIR conference on research and development in information retrieval* (pp. 749–750). ACM.
- Weibull, W. (1951). A statistical distribution function of wide applicability. *Journal of Applied Mechanics – Transactions of the American Society of Mechanical Engineers*, *18*(3), 293–297.
- Williams, B., Pfeiffer, J., & Waller, V. (2013). Lecture capture: student hopes, instructor fears. In *Proceedings of the ASCILITE - Australian Society for Computers in Learning in Tertiary Education Annual Conference* (pp. 913–924).
- Wondmagegnehu, E. T., Navarro, J., & Hernandez, P. J. (2005). Bathtub shaped failure rates from mixtures: A practical point of view. *IEEE Transactions on Reliability*, *54*(2), 270–275.
- Woodward, W. A., & Gunst, R. F. (1987). Using mixtures of Weibull distributions to estimate mixing proportions. *Computational Statistics & Data Analysis*, *5*(3), 163–176.
- Yoon, C., Oates, G., & Sneddon, J. (2014). Undergraduate mathematics students’ reasons for attending live lectures when recordings are available. *International Journal of Mathematical Education in Science and Technology*, *45*(2), 227–240.