Journal of Data Science 21 (1), 43–56
January 2023
DOI: 10.6339/22-JDS1055
Data Science in Action

Motion Picture Editing as a Hawkes Process

Nick Redfern[1]
[1]*Independent researcher, United Kingdom*

Abstract

In this article I analyse motion picture editing as a point process to explore the temporal structure in the timings of cuts in motion pictures, modelling the editing in 134 Hollywood films released between 1935 and 2005 as a Hawkes process with an exponential kernel. The results show that the editing in Hollywood films can be modelled as a Hawkes process and that the conditional intensity function provides a direct description of the instantaneous cutting rate of a film, revealing the structure of a film's editing at a range of scales. The parameters of the exponential kernel show a clear trend over time to a more rapid editing style with an increase in the rate of exogenous events and small increase in the rate of endogenous events. This is consistent with the shift from a classical to an intensified continuity editing style. There are, however, few differences between genres indicating the consistency of editing practices in Hollywood cinema over time and different types of films.

Keywords *computational film analysis; film editing; Hollywood cinema; point process; time series analysis*

1 Introduction

Motion picture editors shape viewers' experience of time in the cinema by controlling the timing of shots by deciding which frame to cut on, where to place a shot within a sequence, how long a shot is held on screen, and the pacing of a film by controlling the rate at which cuts occur and information is presented to the viewer (Pearlman, 2017). The rate at which cuts occur in motion pictures has been shown to affect viewers' levels of arousal and attention (Lang et al., 1999; Ludwig and Bertling, 2017), memory (Lang et al., 2000), and temporal perception and judgements (Balzarotti et al., 2021).

In this article I am interested in motion picture cutting rates as an element of cinematic pacing and I focus on the times at which the cuts in a film occur rather than the durations of the shots themselves. It therefore makes sense to think of editing as a point process. A point process is a stochastic process over a time interval $(0, T]$, whose realisations are the times at which the events comprising that process occur. A one-dimensional temporal point process can be represented simply as a timeline with the timings of events marked on that line. The counting process ($N(t)$) is the number of events (N) that have occurred by time t and is a simple non-decreasing function that increases by 1 each time an event occurs.

For editing in a motion picture, the point process is the set of times (T) at which cuts occur: $T = t_i, \ldots, t_N$, where t_i is the time of the i-th cut and N is the total number of cuts. t_N is taken as the end of the final shot of a film, where a shot is defined as the time elapsed between two

* Email: nickredfernres@outlook.com.

Received April 1, 2022; Accepted June 20, 2022

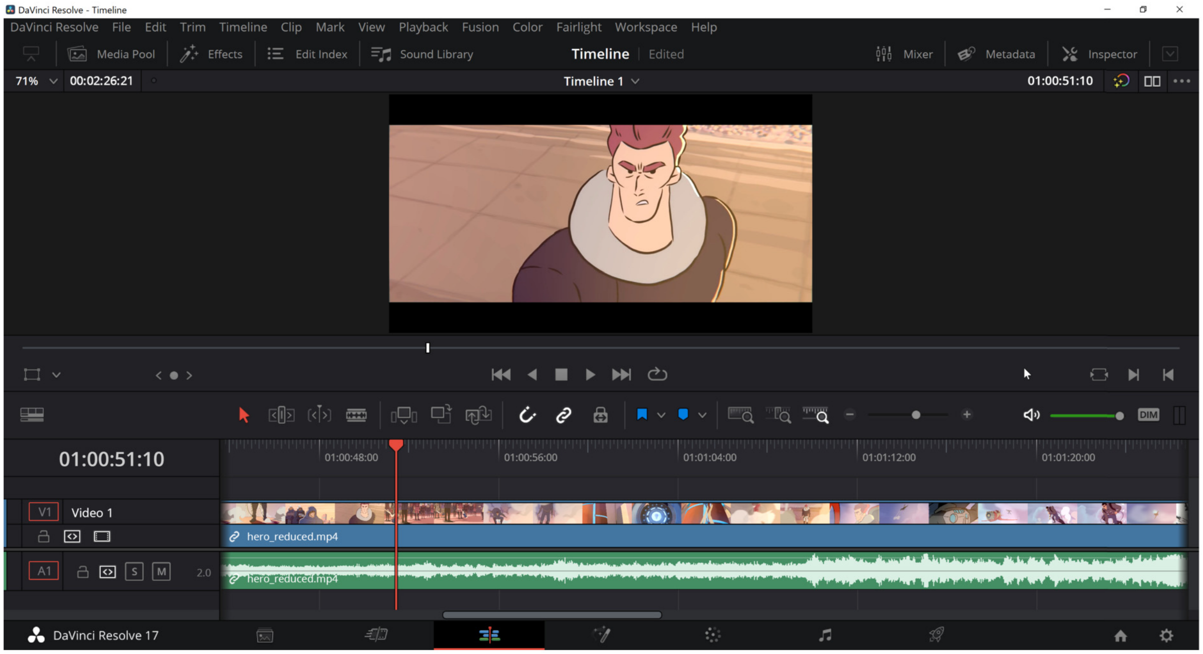

Figure 1: Timeline-based editing in DaVinci Resolve. Clips are arranged in the timeline chronological order.

Table 1: Extract from an edit decision list (EDL) for *The Butcher Boy* (1917) produced by DaVinci Resolve. Each edit in the EDL appears as a numbered event that contains the shot number, reel number, edit type (V = video, C = cut), source timecode in and out points (hours:minutes:seconds:frames), and record timecode in and out points.

Shot	Reel	V	C	Source In	Source Out	Record In	Record Out
001	001	V	C	00:00:00:00	00:00:08:10	00:00:00:00	00:00:08:10
002	001	V	C	00:00:08:10	00:00:13:01	00:00:08:10	00:00:13:01
003	001	V	C	00:00:13:01	00:00:27:05	00:00:13:01	00:00:27:05
004	001	V	C	00:00:27:05	00:00:32:20	00:00:27:05	00:00:32:20
005	001	V	C	00:00:32:20	00:00:38:19	00:00:32:20	00:00:38:19
006	001	V	C	00:00:38:19	00:00:44:20	00:00:38:19	00:00:44:20
...	...	...	...	...	...	...	...

cuts: $s_i = t_i - t_{i-1}$. The counting process is simply the number of cuts that have occurred to time t. It is natural to think of editing as a point process because this reflects how films are edited. Non-linear editing systems such as DaVinci Resolve or Adobe Premiere Pro are timeline-based, with video and audio clips, effects, and transitions arranged in chronological order horizontally across the screen (Figure 1). The ordered timecodes of the cuts in a film can be exported as an edit decision list (EDL) for use in post-production (see Table 1). The timeline visualises and the EDL describes the editing in a film as points on a line and although filmmakers do not describe their work in terms point processes, the marking of points on a line representing the running time of a film is nevertheless a key concept in how editors think about time in the cinema.

The identification of suitable models of pacing in the cinema has applications in the development of data-driven filmmaking technologies, with decisions about the pacing of film being taken out of the hands of filmmakers and given over to artificially intelligent editing systems (see, for example, Leake et al., 2017). A key feature of motion picture editing is that the probability of a cut occurring at any time is not independent of the previous cuts resulting in the clustering of edits (Cutting et al., 2010), with long-range dependence evident in editing patterns (Cutting et al., 2018). This eliminates the Poisson process as an adequate model of motion picture editing due to the requirement that the inter-event times of such a process are statistically independent (Zadeh and Sharda, 2015). The clustering and long-range dependence suggest that editing is a self-exciting process and can be described by point process models with memory.

In this article I analyse the editing in Hollywood cinema as a point process to assess if the pacing of a film can be adequately modelled as a Hawkes process (Hawkes, 1971), a common model for self-exciting point processes. In the next section I provide a brief overview of Hawkes processes. Section three describes the analytical methods used in this article and section four presents the results of modelling the editing in a sample of 134 Hollywood films as a Hawkes process with an exponential kernel.

2 Hawkes Processes

A point process is self-exciting if the occurrence of an event makes future events more likely to occur with a corresponding increase in the conditional intensity function. Formally, a self-exciting point process exhibits positive covariance between collections of events in time: $cov(N_{t_2} - N_{t_1}, N_{t_3} - N_{t_2}) > 0$, for $t_1 < t_2 < t_3$. A self-exciting point process can produce the type of clustering behaviour we see in motion picture editing and so a model that describes this type of behaviour could be a good candidate to describe this aspect of film style.

A widely used class of self-exciting point process models are Hawkes processes, which have been applied to seismology (Ogata, 1988), epidemiology (Unwin et al., 2021), neuroscience (Lambert et al., 2018), finance (Hawkes, 2018), and social media (Watine et al., 2022). A point process can be completely described by its conditional intensity function, $\lambda(t)$, which is the instantaneous rate of events at time t:

$$\lambda(t) = \lim_{\Delta t \to 0} \frac{E\{N(t + \Delta t)\,|H_t\}}{\Delta t},$$

where H_t is the history of the process up to and including time t. A point process is said to be a Hawkes process if the conditional intensity function takes the form

$$\lambda(t|H_t) = \lambda_0 + \sum_{t_i < t} h(t - t_i) \ ,$$

where λ_0 is the background rate at which exogenous events occur, and h is a kernel function that determines how past events influence the intensity at time t and represents the memory of the process. The effects of past events are cumulative so that the whole history of the process affects the conditional intensity function at t but those effects decay over time so that the most recent events exert the greatest influence. The kernel function can be decomposed into $h = \mu h^*$. $\mu < 1$ is the reproduction mean, which describes the expected number of events triggered once an event occurs and is a measure of the endogeneity of a process. h^* is the reproduction kernel, a normalised density function with $\int h^* = 1$, describing how the influence of past events on t_i decays over time. There are various kernel functions that can be used to model a Hawkes process. A common choice is the exponential kernel, which has the reproduction kernel $h^*(t) = \beta e^{-\beta(t-t_i)} 1_{\{t>0\}}$, where β is the decay parameter of an exponential distribution and $1_{\{t>0\}}$ is the Heaviside step function (Cheysson and Lang, 2021; Filimonov and Sornette, 2015). When an event occurs, the conditional intensity function of a Hawkes process with an exponential kernel increases by $\alpha = \mu\beta$.

3 Methods

3.1 Data

The data set used in this study comprises the shot length data for 134 Hollywood films released between 1935 and 2005 divided into five genres (action, adventure, animation, comedy, and drama) from Cutting et al. (2010), and is available through the Cinemetrics database (www.cinemetrics.lv/database).

3.2 Estimation

I use the *hawkesbow* package (Cheysson, 2021) for the statistical programming language R (R Core Team, 2021) to fit Hawkes models with an exponential kernel for each film in the sample. *hawkesbow* uses the function `nloptr()` from the package *nloptr* (Johnson, n.d.), employing the NLOPT-LD-LBFGS algorithm (Liu and Nocedal, 1989) to maximise the likelihood function

$$\mathcal{L}_n(\theta) = \left(\prod_{i=1}^{N} \lambda(t_i)\right) \exp\left(-\int_0^T \lambda(s)\, ds\right),$$

for an observed set of events at times $t_i, \ \ldots, \ t_N$ in an interval $[0, T]$ over the set of parameters θ. The tuple of initial values used when fitting the exponential kernel is (0.2, 0.2, 0.2).

Attempts to fit a Hawkes model with a power law kernel to the Hollywood shot length data proved unsuccessful, with a high number of failures across a range of initial values and inconsistent results between runs, leading to the model being both rejected and not rejected when running the algorithm multiple times against the same data with the same initial values.

3.3 Goodness-Of-Fit

According to the time-rescaling theorem, if the model fitted to the data is correct then the event times of a point process can be rescaled as a Poisson process with a unit rate (Daley and Vere-Jones, 2003, p. 261). For a counting process $N(.)$, the non-decreasing function

$$\Lambda(t) = \int_0^t \lambda(s)\, ds$$

is the *compensator* of the counting process. The transformed sequence of time points $\{t_1^*, t_2^*, \ldots\} = \{\Lambda(t_1), \ \Lambda(t_2), \ \ldots\}$ is a Poisson process with a unit rate if the original series of time points $\{t_1, \ t_2, \ \ldots\}$ is a realisation of a point process defined by the compensator $\Lambda(.)$.

Goodness-of-fit can then be assessed by a two-sided Kolmogorov-Smirnov test of the null hypothesis that the t_i^*/t_N^*, $1 \leqslant i \leqslant N$, are distributed as order statistics from a standard uniform distribution ($U_{[0,\ 1]}$) (Laub and Taimre T, 2021, pp. 79-84). If the model is a good fit, a plot of the cumulative distribution function of the rescaled times should lie along a 45-degree line within the 99.9% confidence band defined by $1.95/\sqrt{N}$, for a Kolmogorov-Smirnov test with $\alpha = 0.001$. Applying a Bonferroni correction for multiple testing with 134 films in the sample results in a family-wise error rate of $0.001 \times 134 = 0.134$.

Due to limitations of space, I include only one example of an individual film here. The complete set of plots for all the films in the sample is available in the supplementary material for this article.

4 Editing in Hollywood Cinema as a Hawkes Process

The results of fitting the exponential kernel show that motion picture editing can be modelled as a Hawkes process, with the model well-fitting for all 134 films in the sample.

To illustrate the results of fitting the Hawkes model to the editing of an individual film, Figure 2 plots the results of modelling the time series data for *The Perfect Storm* (2000), which tells the story of the crew of the fishing boat Andrea Gail that was lost at sea in the 'perfect storm' at the confluence of two powerful weather fronts and a hurricane in October 1991. Figure 2.A plots the counting process and Figure 2.B plots the goodness-of-fit of the rescaled times against a standard uniform distribution along with the results of the Kolmogorov-Smirnov test. The counting process shows how the cutting rate changes over time, with an accelerating cutting rate until the final fifth of the film, when the rate at which cuts occur slows down. It is the conditional intensity function in Figure 2.C that is most informative, describing how the cutting rate changes over the course of the film and enabling the analyst to identify potentially explicable features at different scales.

For example, the conditional intensity function shows an increasing trend until 5790.7s (~1:36:30), which takes us to the point when the helicopter rescue crew is forced to abandon its search for the boat and ditches in the water, while the crew of the Andrea Gail loses hope and Captain Billy Tyne admits, 'Boys, that's it. We can't make it. We're turning around.' After this point the cutting rate decreases as the ship is finally overwhelmed by a rogue wave and their families comes to terms with the loss of the crew. This suggests that the cutting rate at this scale is associated with emotion rather than action, tied to the hopes and fears of the characters rather than the effects of the storm itself. Within the large-scale evolution of the conditional intensity function there are mid-scale features at the level of the sequence. The conditional intensity function shows a downward trend from 3046.2-4239.7 seconds (~0:50:46-1:10:43). During this sequence we see the crew of the Andrea Gail make the fateful decision to head further out to sea where they land a large catch before being forced to return home due to a broken ice-machine. As the crew's fortunes improve the cutting rate decreases, again linking the editing to the dominant emotion of a sequence, as the crew's pessimism about their lack of success is replaced by their optimism about the future. Within this sequence are small-scale features in which the conditional intensity function rises sharply before dropping away equally quickly to a floor consistent with the trend of the cutting rate at the scale of the sequence. These local peaks are associated with moments of intense action when Murph, one of the fishermen, falls overboard and is rescued (3218s, ~0:53:38), and a yacht thrown about by the storm calls for assistance (3705.9, ~1:01:46). These local peaks inject moments of danger into a sequence with an overall editing trend associated with the positive emotions of hope and optimism. These results indicate that the editing of *The Perfect Storm* serves different functions at different scales.

Figure 3 plots the trends in the model parameters for the whole corpus of films over time. There is a significant increase in the background rate λ_0 from 1955 to 2005 measured by Spear-

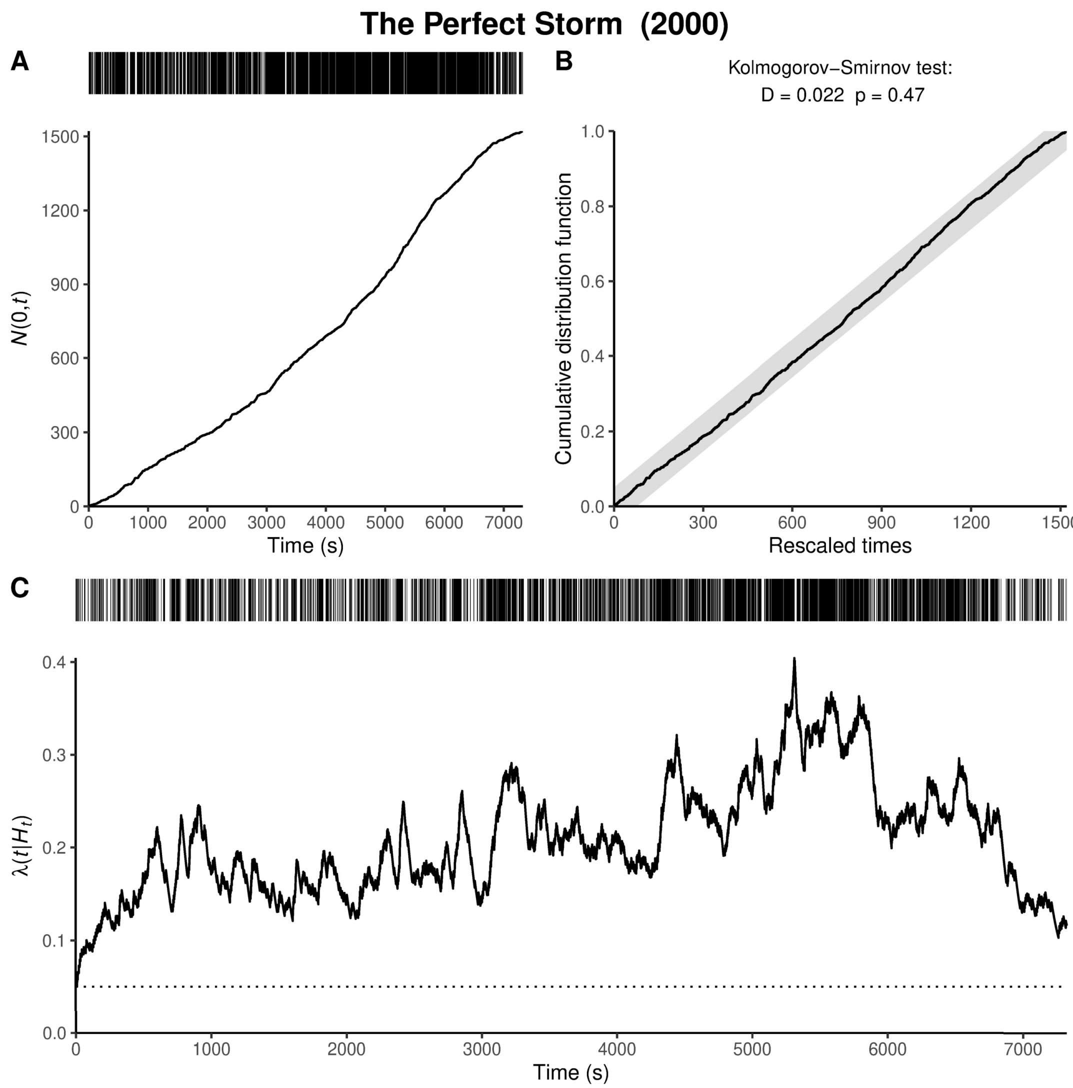

Figure 2: Editing in *The Perfect Storm* (2000) modelled as a Hawkes process with an exponential kernel ($\lambda_0 = 0.0502$, $\mu = 0.763$, and $\beta = 0.0106$). (A) Counting process. (B) K-S plot: the rescaled times are plotted against the cumulative distribution function of a standard uniform distribution ($U_{[0,\ 1]}$), with a 99.9% confidence band. (C) Conditional intensity function of the fitted model.

man's rho ($r_s(132) = .53$ ($95\%CI$: .40, .64), $p \leqslant .001$), with an intensifying trend after 1975. This trend tracks the shift to an intensified continuity style in Hollywood cinema after 1975 characterised by increased cutting rates, with editors cutting on every line of dialogue and including more reaction shots compared to editing practices in classical sound era (1930-1960) (Bordwell, 2006; Cutting and Candan, 2015; Redfern, 2020). The range of the reproduction mean in the sample is broad, from 24% in *Harvey* (1950) to 86% for *Ghost* (1990), but there is only a weak increasing trend over time from 1935 to 2005 ($r_s(132) = .22$ ($95\%CI$: .06, .39), $p = .01$), indicating that while the rate at which exogenous events occurs has increased, the endogeneity of the process has changed little over time. Films released before 1960 have a slightly higher repro-

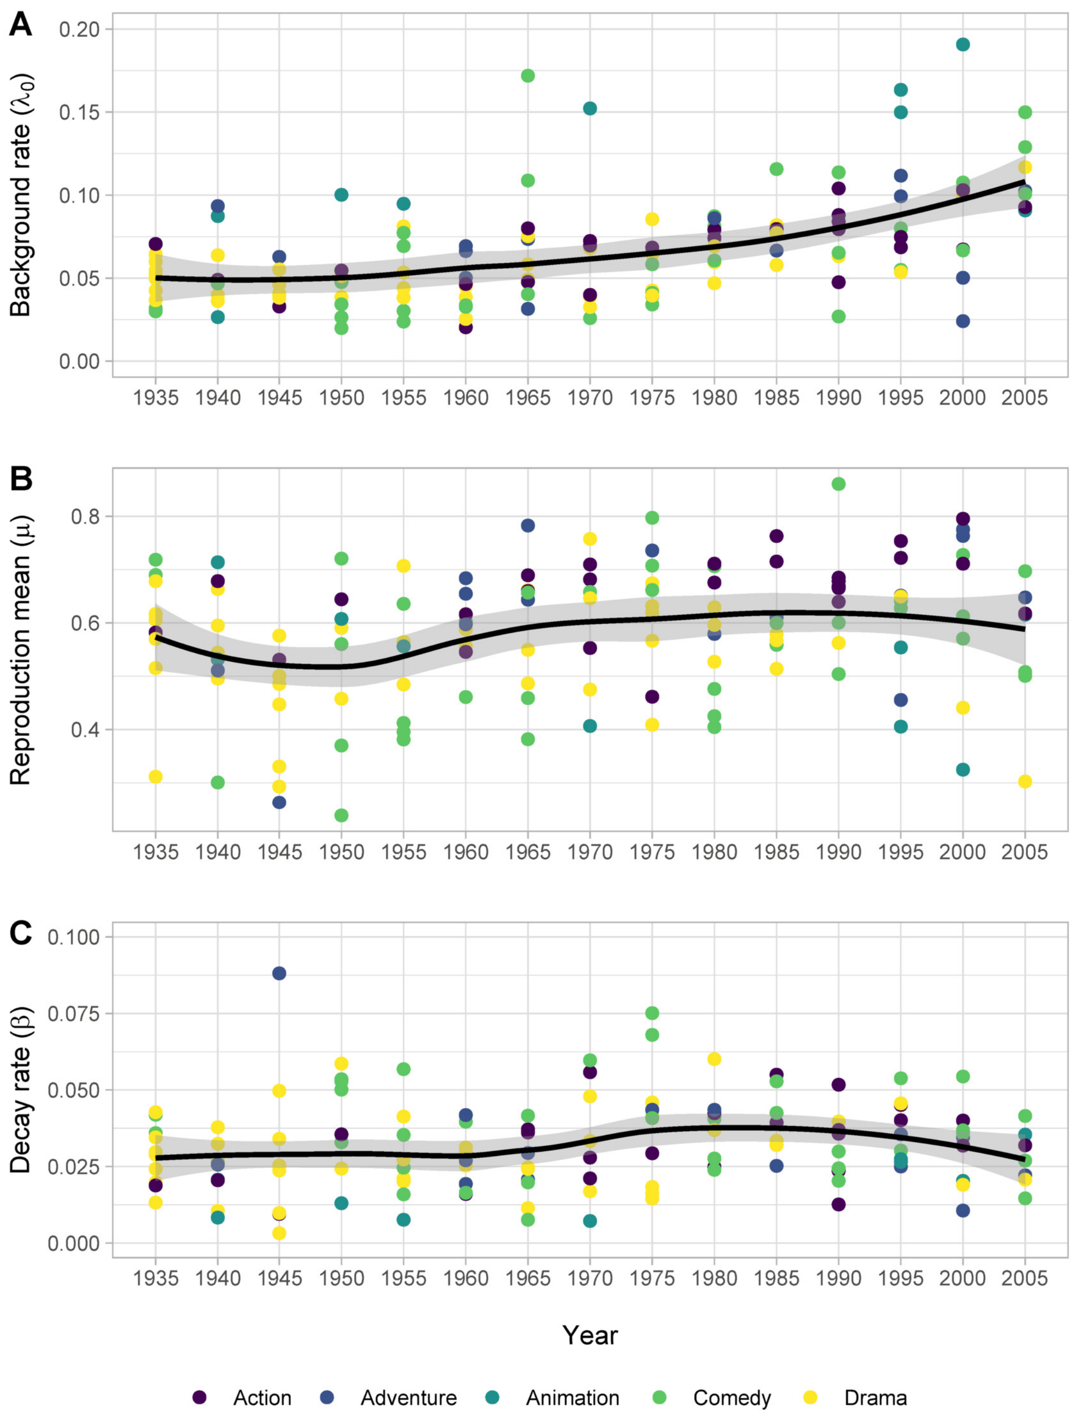

Figure 3: Trends in the (A) background rate (λ_0), (B) reproduction mean (μ), and (C) decay rate (β) of a Hawkes process with an exponential kernel fitted to shot length data of 134 Hollywood films released from 1935 to 2005, with lowess (locally weighted scatterplot smoothing) trendlines and 95% confidence interval. All lowess trendlines are have degree = 2 with a span of 0.75.

duction mean than those released from 1965, which again reflects the shift from a classical to post-classical mode of filmmaking. There is no evidence that the trend in μ increased after 1975 similar to that of the background rate, and the reproduction mean has remained constant since 1965 with an average of approximately 60% of cuts in films in the sample occurring as a response to an exogenous cut. There is a small decline in the value of μ for films released in 1945, though this is due to a small number of films (*Detour*, *In Pursuit to Algiers*, *The Lost Weekend*) released in this year that have atypically low values of μ. The decay rate of the reproduction kernel shows no trend over time ($r_s(132) = .16$ $(95\%CI : -.01, .32)$, $p = .06$) but increased slightly in the 1970s and 1980s. Again, this was a temporary change in style and from 1990 the trend in the decay rate decreases so that by 2005 it is comparable with that of 1935. A reproduction kernel with a higher decay rate has shorter memory as the density of the exponential distribution falls away more quickly as β becomes larger. This indicates that the influence of cuts in films of the new Hollywood era did not extend as far in time as in other eras, resulting in clusters of shots of similar size in the 1970s (as reflected in the constant trend of the reproduction mean) but which were more loosely constructed than in other eras. However, it should again be noted that the change in style is small and may represent a tendency for drama films to be under-represented in the sample during these years rather than a clear change in style. The overall pattern is one of stylistic consistency in Hollywood filmmaking as the underlying principles of continuity editing have persisted over time and the relationships between cuts have been maintained even though cutting rates have evolved over time.

Figure 4 plots λ_0, α, and β for the Hollywood films in the sample and again shows the stylistic consistency across the 70-year span of the data set, with all but one film (*In Pursuit to Algiers*) lying close to a plane in a 3-D space defined by the linear model $\beta = 0.018\lambda_0 + 1.315\alpha + 0.006$. Drama films tend to lie slightly above the plane while action films tend to lie slightly below due to the difference in their reproduction means. Simulated Hawkes processes for creating edited film sequences can be generated by selecting the parameters (λ_0, $\alpha = \mu\beta$, β) that define a location on or near to this plane.

Figure 5 plots the distributions of the parameters of the fitted exponential kernels for each genre in the sample. To compare the distribution of the parameters across genres I performed a Kruskal-Wallis test, with pairwise post-hoc Wilcoxon-Mann-Whitney tests assuming an family-wise error rate of .10 and 10 tests giving a Bonferroni-adjusted p-value of .01. The effect size is given by Cliff's d (Cliff, 1993). There is a significant difference for the background rates ($\chi^2(4) = 17.52$, $p = < .01$), with significant pairwise differences between the animation films and the action ($p =< .01$, $d = 0.69$) and drama ($p =< .01$, $d = 0.75$) genres. There is a significant difference among genres for the reproduction mean ($\chi^2(4) = 23.01$, $p = < .01$), with values of μ for the action genre significant different from those of animation ($p =< .01$, $d = 0.65$), comedy ($p =< .01$, $d = 0.42$), and drama ($p =< .01$, $d = 0.65$) films. There are no other significant pairwise differences among genres. The omnibus test for the decay rate also shows a significant difference ($\chi^2(4) = 14.46$, $p = < .01$), with pairwise differences between the animation genre and the action ($p =< .01$, $d = -0.65$) and comedy ($p =< .01$, $d = -0.69$) genres.

There is no consistent pattern of differences between genres, with only the animated films exhibiting large differences from some other genres. The background rate for these films tends to be higher than those of some other genres, while the decay rate tends to be lower. From Figure 5 this genre is also the only group of films with apparent sub-groups, marking a difference between those films for which λ_0 and μ are higher than other films (*The Aristocats* (1970), *Pocahontas* (1995), *Toy Story* (1995), *Dinosaur* (2000)) and those for which the opposite is the

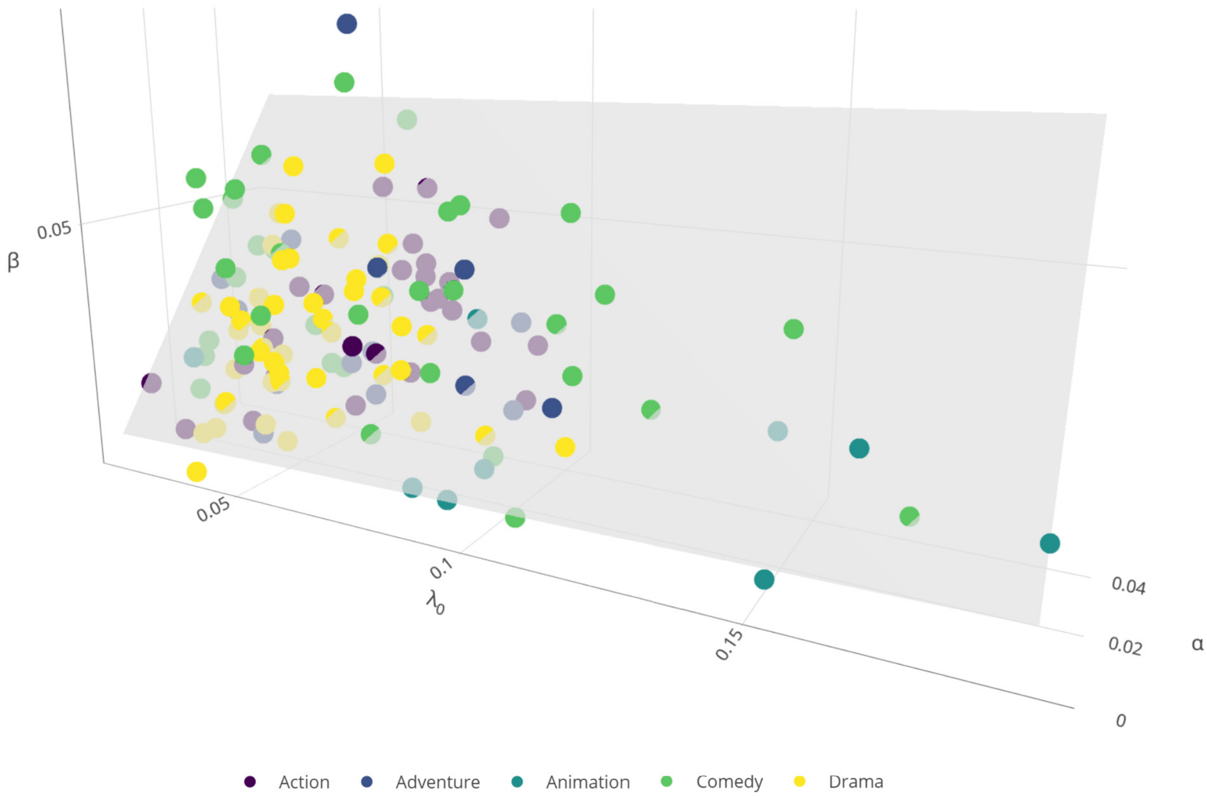

Figure 4: The parameters λ_0, $\alpha = \mu\beta$, and β of a Hawkes process with exponential kernel fitted to 134 Hollywood films released from 1935 to 2005, with fitted 2D-plane fitted by the linear model $\beta = 0.018\lambda_0 + 1.315\alpha + 0.006$. An interactive version of this plot is available as part of the supplementary material for this article.

case (*Pinocchio* (1940), *Cinderella* (1950), *Lady and the Tramp* (1955), *Madagascar* (2005)). *Fantasia* (1940) belongs to the latter group but is distant from the other animated films, with a notably lower background rate ($\lambda_0 = 0.0265$) and higher reproduction mean ($\mu = 0.7136$). The values of the decay rate do not show the same distinction between these groups, though films released after 1970 tend to have higher values of β than those released earlier. Except for *Madagascar*, which was released in 2005, this indicates a difference in the pacing of animated films prior to 1970 and those released later, with the latter group having both higher cutting rates and but with cuts that have less influence on the timing of subsequent edits.

5 Conclusion

This article is the first to consider if pacing in the cinema can be modelled as a Hawkes process. The results show that the conditional intensity function of a Hawkes model with an exponential kernel can provide a good estimate of the cutting rate of Hollywood films and is informative at different levels of analysis. At the level of the individual film, the conditional intensity function is a direct way of representing the instantaneous cutting rate of a film that has descriptive power, straightforwardly communicating how the editing of a film evolves over the course of its running time, and analytical power, supporting a bottom-up approach to analysing film style

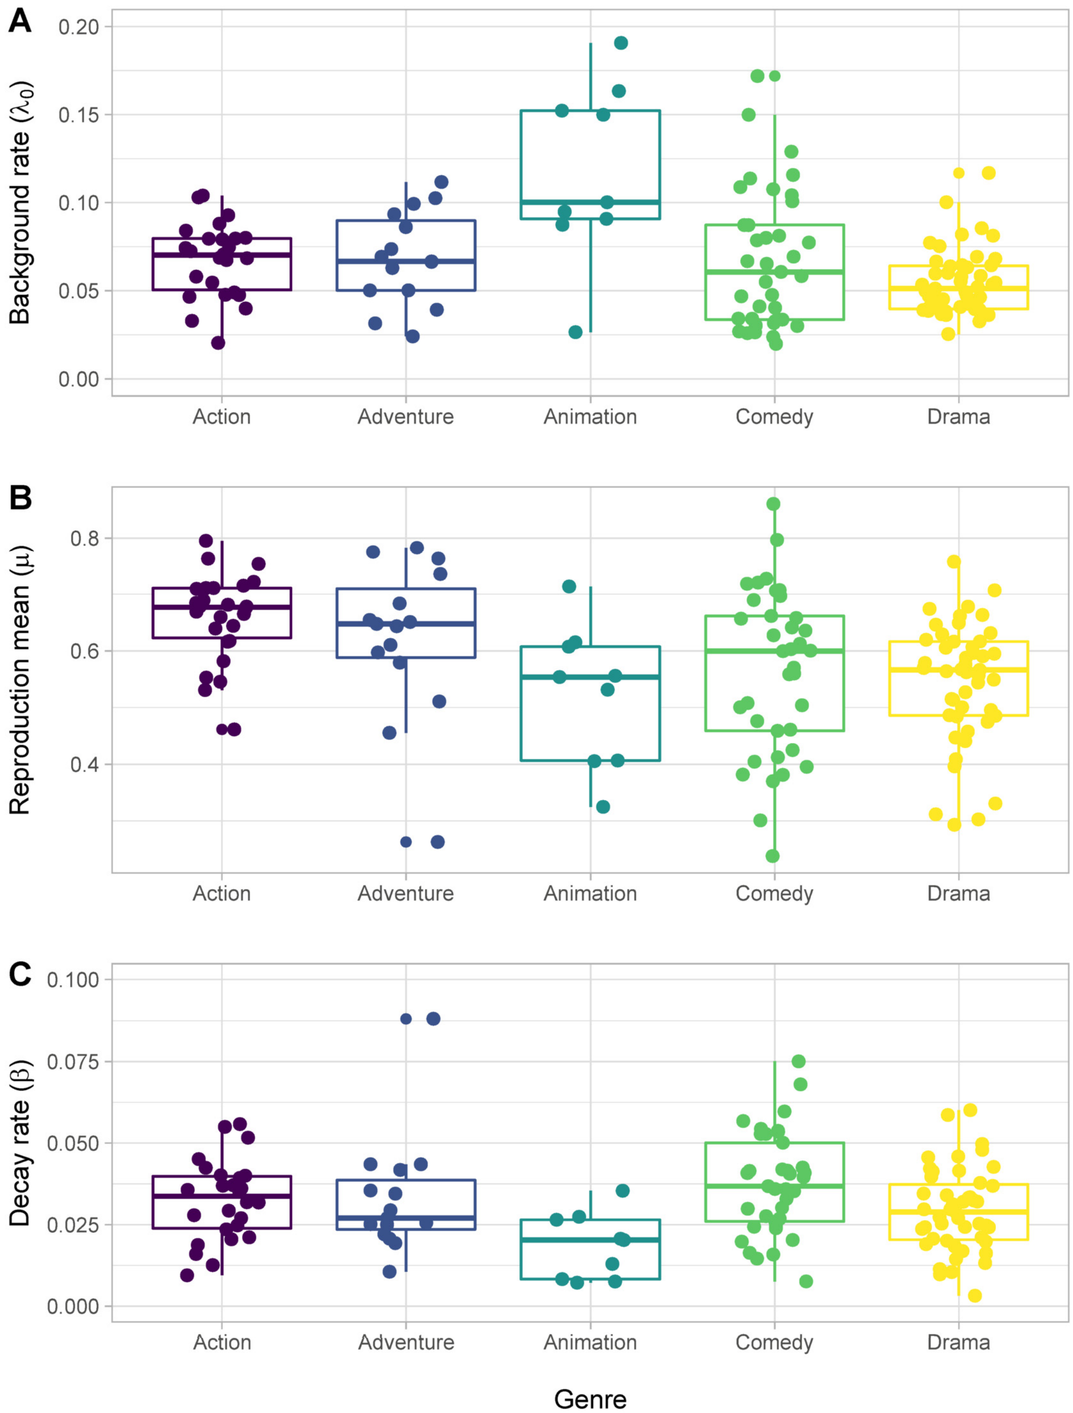

Figure 5: The (A) background rate (λ_0), (B) reproduction mean (μ), and (C) decay rate (β) of a Hawkes process with an exponential kernel fitted to shot length data of 134 Hollywood films released from 1935 to 2005 by genre.

that allows the researcher to identify potentially explicable features at the different scales of the scene, sequence, and film (Redfern, 2015). At the level of the corpus, the parameters of the model allow us to identify stylistic trends over time. The shift to a more rapid editing style is evident in the increase in the background rate, but there is only a small change in the reproduction mean over time and no trend in the reproduction kernel over 70 years of filmmaking. There is no evidence of shift to a more rapid cutting style associated with the arrival of an 'MTV aesthetic' in the early 1980s, which apparently saw filmmakers adopt an aesthetic derived from music videos (Calavita, 2007). On the contrary, the emergence of intensified continuity of the post-classical Hollywood era would support the conclusion that while the pacing of Hollywood films has evolved, the continuity style of editing films has remained largely stable with aesthetic norms persisting over time (Bordwell, 2002). The trend in the background rate clearly indicates this change began long-before the advent of MTV in the 1960s (see also Redfern, 2014). Action and animated films exhibit some differences from other genres, but in general there are no differences in the model parameters between genres.

The results presented here have some limitations to be explored in future research. First, I assumed that the background rate of the Hawkes process is constant. Although the results presented here indicate that a Hawkes model with a constant background rate can model the editing of Hollywood films adequately, this may not be a realistic assumption. Cutting (2016a,b) notes that the style of Hollywood films evolves over the course of their running time, with different editing regimes associated with different stages of a film. Better results may therefore be achieved using a Hawkes model with a time-varying background rate. As the four-act structure is common to Hollywood filmmaking (Thompson, 1999), piecewise linear or log-linear baselines with four knots or basis functions (Omi et al., 2017) may provide a better fitting model to deal with non-stationary features. Second, in this article I have assumed that the parameters of the exponential kernel are a reasonable model for the memory of editing as a point process. Again, the extent to which this assumption is realistic requires further consideration and non-parametric estimation of the Hawkes process parameters may better fit the data. Third, estimation of the coefficients of the regression plane in Figure 4 could be improved by taking into account the increasing trend over time exhibited by the background rate λ_0 to include release year as a variable in the linear model. Editing practices in Hollywood cinema comprise a set of norms that are largely consistent across different types of films, but fitting regression planes for separate genres may identify differences between different types of films. Including these variables would allow for the simulation of Hawkes processes that could vary according to the genre of a film or capture different aesthetic norms in Hollywood editing, such as the distinction between classical and post-classical continuity editing (Bordwell, 2002).

Fourth, I have considered editing as a univariate point process in isolation from other elements of film style. Analysing the temporal structure of a film as marked Hawkes process will allow for other elements of film style, such shot scale, camera movement, camera angle, or any state of the narrative we care to define (e.g., dialogue, action, complication, climax, etc.), to be incorporated into the model. Potentially interesting models are Gaussian Marked Hawkes Process, which combine a Hawkes process to model events in continuous time with a Gaussian distribution for modelling the meta information assigned to events (Seonwoo et al., 2018); Markov-modulated Hawkes processes (Wang et al., 2012), which allow for a Hawkes process to switch between states; and neural Hawkes processes (Mei and Eisner, 2017), which combine a Hawkes process with a recurrent neural network to better capture the real world behaviour of self-exciting point processes and which allows for past events to either excite *or* inhibit future events. Given the complexity of editing as a temporal process, in which the cutting rate of a film is determined by its genre, the stage of the film at which a cut occurs, and the type of shot selected, flexible multivariate models will capture key features overlooked when using a univariate model.

Finally, this article only considered the case of editing in Hollywood cinema and other modes of filmmaking with different sets of aesthetic norms, such as art cinema, documentary filmmaking, or television production, may not be well-fitted by a Hawkes model with an exponential kernel. They may not exhibit self-exciting behaviour at all. Future research will therefore need to examine different types of filmmaking as a point process.

Supplementary Material

The complete set of results and plots for all 134 films in the sample along with the R code used in this project are available for the reader to explore as a shiny app at https://tinyurl.com/2p8c86u3. The data, code, and results for this article are also available on the supporting GitHub repository at DrNickRedfern/hollywood-hawkes.

References

Balzarotti S, Cavaletti F, D'Aloia A, Colombo B, Cardani E, Ciceri MR, et al. (2021). The editing density of moving images influences viewers' time perception: the mediating role of eye movements. *Cognitive Science*, 45(4): e12969.

Bordwell D (2002). Intensified continuity visual style in contemporary American film. *Film Quarterly*, 55(3): 16–28.

Bordwell D (2006). *The Way Hollywood Tells It: Story and Style in Modern Movies*. University of California Press, Berkeley, CA.

Calavita M (2007). "MTV Aesthetics" at the movies: interrogating a film criticism fallacy. *Journal of Film and Video*, 59(3): 15–31.

Cheysson F (2021). hawkesbow: Estimation of Hawkes Processes from Binned Observations. R package version 1.0.2.

Cheysson F, Lang G (2021). *Strong-mixing rates for Hawkes processes and application to Whittle estimation from count data*.

Cliff N (1993). Dominance statistics: ordinal analyses to answer ordinal questions. *Psychological Bulletin*, 114(3): 494–509.

Cutting JE (2016a). The evolution of pace in popular movies. *Cognitive Research: Principles and Implications*, 1(1): 30.

Cutting JE (2016b). Narrative theory and the dynamics of popular movies. *Psychonomic Bulletin & Review*, 23(6): 1713–1743.

Cutting JE, Candan A (2015). Shot durations, shot classes, and the increased pace of popular movies. *Projections*, 9(2): 40–62.

Cutting JE, DeLong JE, Brunick KL (2018). Temporal fractals in movies and mind. *Cognitive Research: Principles and Implications*, 3(1): 8.

Cutting JE, DeLong JE, Nothelfer CE (2010). Attention and the evolution of Hollywood film. *Psychological Science*, 21(3): 432–439.

Daley D, Vere-Jones D (2003). *An Introduction to the Theory of Point Processes: Volume I: Elementary Theory and Methods*. Springer International Publishing, New York, NY.

Filimonov V, Sornette D (2015). Apparent criticality and calibration issues in the hawkes self-excited point process model: application to high-frequency financial data. *Quantitative Finance*, 15(8): 1293–1314.

Hawkes AG (1971). Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1): 83–90.

Hawkes AG (2018). Hawkes processes and their applications to finance: a review. *Quantitative Finance*, 18(2): 193–198.

Johnson SG (n.d.) The NLopt nonlinear-optimization package.

Lambert RC, Tuleau-Malot C, Bessaih T, Rivoirard V, Bouret Y, Leresche N, et al. (2018). Reconstructing the functional connectivity of multiple spike trains using Hawkes models. *Journal of Neuroscience Methods*, 297: 9–21.

Lang A, Bolls P, Potter RF, Kawahara K (1999). The effects of production pacing and arousing content on the information processing of television messages. *Journal of Broadcasting & Electronic Media*, 43(4): 451–475.

Lang A, Zhou S, Schwartz N, Bolls PD, Potter RF (2000). The effects of edits on arousal, attention, and memory for television messages: when an edit is an edit can an edit be too much? *Journal of Broadcasting & Electronic Media*, 44(1): 94–109.

Laub PJ, Taimre T LY (2021). *The Elements of Hawkes Processes*. Springer International Publishing, Cham.

Leake M, Davis A, Truong A, Agrawala M (2017). Computational video editing for dialogue-driven scenes. *ACM Transactions on Graphics*, 36(4): 1–14.

Liu DC, Nocedal J (1989). On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45(1–3): 503–528.

Ludwig M, Bertling C (2017). The effect of cutting rates on the liking of live sports broadcasts. *International Journal of Sport Communication*, 10(3): 359–370.

Mei H, Eisner JM (2017). The neural hawkes process: a neurally self-modulating multivariate point process. In: *Advances in Neural Information Processing Systems* (I Guyon, UV Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan, R Garnett, eds.), volume 30. Curran Associates Inc.

Ogata Y (1988). Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical Association*, 83(401): 9–27.

Omi T, Hirata Y, Aihara K (2017). Hawkes process model with a time-dependent background rate and its application to high-frequency financial data. *Physical Review E*, 96: 012303.

Pearlman K (2017). Editing and cognition beyond continuity. *Projections*, 11(2).

R Core Team (2021). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Redfern N (2014). Robust estimation of the mAR index of high grossing films at the US box office, 1935 to 2005. *Journal of Data Science*, 12(2): 277–291.

Redfern N (2015). Exploratory data analysis and film form: the editing structure of *Friday the Thirteenth* (1980). *Post Script: Essays in Film and the Humanities*, 34(2–3): 71–83.

Redfern N (2020). Modelling shot lengths of Hollywood motion pictures with the Dagum distribution. *Statistica*, 80: 81–99.

Seonwoo Y, Oh A, Park S (2018). Hierarchical Dirichlet Gaussian marked Hawkes process for narrative reconstruction in continuous time domain. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 3316–3325". Association for Computational Linguistics, Brussels.

Thompson K (1999). *Storytelling in the New Hollywood*. Harvard University Press, Cambridge, MA.

Unwin HJT, Routledge I, Flaxman S, Rizoiu MA, Lai S, Cohen J, et al. (2021). Using Hawkes processes to model imported and local malaria cases in near-elimination settings. *PLOS Computational Biology*, 17(4): e1008830.

Wang T, Bebbington M, Harte D (2012). Markov-modulated Hawkes process with stepwise decay. *Annals of the Institute of Statistical Mathematics*, 64(3): 521–544.

Watine P, Bodaghi A, Schmitt KA (2022). An attempt to evaluate information spread using the Hawkes process. In: *Proceedings of the Future Technologies Conference (FTC) 2021, Volume 2* (K Arai, ed.), volume 359, 354–362. Springer International Publishing, Cham.

Zadeh AH, Sharda R (2015). Hawkes point processes for social media analytics. In: *Reshaping Society through Analytics, Collaboration, and Decision Support* (LS Iyer, DJ Power, eds.), volume 18, 51–66. Springer International Publishing, Cham.