Evolution in Hollywood editing patterns?

Mike Baxter
16 Lady Bay Road, West Bridgford, Nottingham, NG2 5BJ, U.K.
(e-mail: michaelj.baxter@btconnect.com)

November 2013

Abstract *Cutting, Delong and Nothelfer (2010) use statistical methods to investigate the evolution of shot-length patterns in popular film. They argue, using what they call a 'modified autoregressive index' (mAR), that patterns are becoming increasingly clustered and also evolving towards 1/f structure, a pattern described in later publication as 'like those that our minds may naturally generate'. This paper shows that the interpretation of the mAR index is wrong. It is also shown that that the results concerning 1/f patterns can be interpreted in an equally plausible and much less 'exciting' way. That is, although there are undoubtedly interesting temporal patterns in the shot length structure, they can't be interpreted in terms of the 'evolution of Hollywood film' in the sense intended in the original paper.*

1 Introduction

Cinemetrics has been characterized as an online research tool to facilitate the scholarly study of film editing. There is an emphasis on cutting rates as evidenced by shot-lengths (SLs) and the relationship between cutting rates and the history of film is singled out as a potential area of study (Bosse, Tsivian, and Brisson, 2011). It involves the analysis of quantified 'filmic' data using statistical methods that are, perhaps, alien to a majority of film scholars. Many of the ideas currently being explored can be traced back to Barry Salt's 1974 paper *Statistical Style Analysis of Motion Pictures*, and early applications are to be found in Salt's own work, particularly his book *Film Style & Technology: History & Analysis* (2009), now in its third edition but first published in 1983.

Salt's emphasis on quantification as a basis for the objective study of film style has been described, in a review by Abel (1989, p. 48), as 'in the tradition of British empiricism' or what Salt calls 'Scientific Realism'. This has echoes of the quantitative revolutions in Geography and Archaeology that were well underway when Salt first published. These 'revolutions' had a considerable impact on practice; quantifiers were often labeled 'positivists' - as a badge of pride by themselves or as a term of abuse by those convinced that the advocacy of 'scientific' study dehumanized their discipline.

In film studies, by contrast, my impression is that the idea of quantification has met largely with indifference and/or incomprehension rather than enthusiasm or vehemence, and Salt's methodologies were not much emulated. This is suggested as much in the preface to *The Numbers Speak*, an essay published for the first time in Salt (2006, pp. 389-397) but based on work undertaken much earlier that develops, in mathematical language, an idea first broached in Salt (1974). Specifically he notes that the work was based on research undertaken over a long period that had been dropped at one point but which he now thought important enough to publish 'given that there is a new interest from other people in shot length statistics' (Salt, 2006, p. 388).

These 'other people' include the developers of the *Cinemetrics* website. The measurement software on the website enables users to record the SLs for individual films, and a database of several thousand user-contributed analyses now exists. There are several consequences of this development; it is much easier to generate data than was previously the case; there is a lot more of it; and it is 'richer' than was often the case for earlier data collection. The obvious question is what can you usefully do with it, and how?

Historically the average shot length (ASL) has been the most widely used quantitative measure of film style. This is simple to understand and, importantly, comparatively simple to measure. All you need to know is the duration of a film and the number of shots in it - it is not necessary to know the individual SLs. This restricts the kind of analysis which can be done, though these are still informative. There is, for example, a nice graph in Salt (2009, p. 378) showing the year-by-year variation in (mean) ASLs for 7792 American films between 1930 and 2006. It shows that the comparatively slow cutting rate associated with the advent of sound in the late 1920s increased fairly rapidly to the mid-1930s but then rose again to a peak (1947-1955) after which there was a steady decline to about 1975, some flattening out for the next 10 years, then a further steady decline.

This example has been chosen to illustrate what, as a statistician, I view as the major benefit of quantification and the availability of large data sets, and that is pattern recognition. Once pattern is recognized it invites explanation and that is the province of the film scholar rather than the statistician. The statistical analysis of quantified filmic data is a means rather than an end but has the potential for revealing non-obvious pattern that demands explanation.

That measuring individual SLs allows 'richer' analysis is simply because much more can be done than computing an ASL. For example it becomes possible to compute the median shot length. It has been argued that this rather than the ASL is a better measure of 'film style' that should be used in preference to the ASL (Redfern, 2012a). Others have questioned this, but the point is that a debate is possible because data are now available.

Although, from a purely statistical point of view, many applications to quantified filmic data involve fairly simple methodology, more complex pattern seeking methods have recently been explored in a series of discussion papers on the *Cinemetrics* website. Many of these involve both the characterization of SL structures within films and their comparison across films.

One of the most interesting and ambitious papers using cinemetrics analysis to have been published is that of Cutting, DeLong and Nothelfer (2010) on the evolution of Hollywood film. It appears to have everything that might be wished for in a cinemetric study. The statistical methodology used is complex - not a merit in itself but needed for the purpose intended. The methodology appears to identify non-obvious patterning in the data for which a filmic interpretation is possible and psychological theories of attention can be advanced by way of explanation. The paper has attracted attention beyond that normally enjoyed by a scholarly publication.

As well as interesting I think the paper is also potentially important, but there is a problem with the interpretation of one of the statistics used that undermines the conclusions drawn from it. It is additionally possible to reinterpret other results in the paper in a way that leads to rather less 'eye-catching' conclusions than those reported.

This paper is an attempt to explain what the statistical issues are and how they affect the conclusions that have been drawn, and subsequently repeated in the literature. There are interesting patterns in the data, but these can be more simply explained than hitherto and raise other questions about the evolution of Hollywood film.

2 Evolution in Hollywood film

Cutting, DeLong and Nothelfer (2010) (henceforth CDN) develop an index of SL patterning that they claim, over time, shows that shots in films have become increasingly clustered into packets of shots of similar length. This is described as 'evolution'; in their words 'film editors and directors have incrementally increased their control over the visual momentum of their narratives, making the relations among shot lengths more coherent over a 70-year period'.

The analysis is linked with claims that 'the shot structure in films has been evolving towards $1/f$ spectra'. This last claim is repeated and elaborated on in later publications using less technical language. Thus Delong *et al.* (in press) state that the earlier analyses 'lead us to believe that film may be evolving; the characteristics of film may change over time to better serve cognitive mechanisms like attention'. Smith *et al.* (2012, p. 109) explain this as meaning that 'whole films are characterized by rhythmic fluctuations that appear to guide viewer attention' and that 'the shot-duration patterns of film might increasingly be like those our minds may naturally generate'. The statement in the same paper that since about 1960 'the shot lengths of films analyzed as sequences across entire films have increasingly approached a $1/f$ pattern' is not qualified in any way.

These claims have attracted attention beyond the academic sphere. *The New Scientist* ran an article entitled *Solved: The Mathematics of the Hollywood Blockbuster* and a web search will bring up several similarly entitled items. CDN do not make such exaggerated claims. Nevertheless, Cutting's remark, in the context of a debate with Barry Salt on the *Cinemetrics* 'Discussion Board', 'that the news was that there was an increasing trend towards $1/f$ over time. Insofar as we know, there is no trend from near randomness towards $1/f$ that has been traced in any domain in any other science, let alone any other art form' is a striking claim that merits careful critical scrutiny.

The analyses in CDN can be broken down into four stages.

1. The SL data for a film are converted, by statistical methods, into indices measuring some property of interest.
2. Statistical methods are used to identify temporal patterning in the indices.
3. A 'filmic' interpretation is offered for the patterns detected, namely they are evolving.
4. Psychological theories of attention are offered to 'explain' this evolution.

In this progression from data collection, through processing, statistical analysis, pattern recognition and interpretation in substantive (filmic and psychological) terms it is inevitably the higher-level and novel substantive claims that attract attention. They rest, however, on the statistical foundations at stages 1 and 2. Should these be unsound the whole edifice is unsafe. The argument in the present paper is that the foundations are unsound. The rest of the paper attempts to explain why.

3 Methodology

A lot of the interest in applied time-series analysis, in other areas of application, lies in predicting future observations from past observations. This is often done in terms of statistical models that postulate a systematic relationship between a current observation and past observations, to which is added a random (or error, or disturbance) term.

There are many ways of specifying the systematic and random components and the way they are combined, but commonly linear models are used. There is also considerable choice among these; one of the simplest, and the only one of concern here – it is the one used in CDN – is the autoregression (AR) model, which assumes that an observation can be predicted from a linear combination of the observed values of previous observations. The order, or index, of the AR model is the number of past terms needed for satisfactory prediction.

Pairs of adjacent observations are separated by a lag of $h = 1$; the autocorrelation coefficient at lag 1, ρ_1, is the correlation between all such pairs; ρ_2 is the correlation between pairs separated by one intervening observation, and so on; $\rho_0 = 1$. The plot of ρ_h against h is the *autocorrelation function* (ACF).

The *partial autocorrelation coefficient* at lag h, α_h, is the correlation between pairs of observations after factoring out the effect of other intervening terms. It can be thought of as measuring the 'predictive power' that still remains after removing the effect of observations closer to the value to be predicted. A variety of algorithms exist for calculating α_h; $\alpha_1 = \rho_1$, and at greater lags the partial autocorrelations can be defined recursively in terms of the autocorrelations, but other methods can be used. The plot of α_h against h for $h \geq 1$ is the *partial autocorrelation function* (PACF).

Faced with real data and a plethora of models that might be used to describe them, choosing an appropriate model is not an easy task. The ACF and PACF are useful diagnostic tools. The PACF is widely used to assess whether an AR model is appropriate. The idea is that so long as observations at lag h (and those with smaller lags) retain predictive power the calculated α_h will be 'significantly' large and will drop sharply to non-significant values once they no longer have predictive value. The lag after which the drop occurs defines the order of the AR model.

Identification of where the 'drop occurs' is frequently not this obvious and statistical aids can be used. In both CDN and Redfern (2012b), who has reanalyzed the CDN data, their initial analyses examine α_h sequentially to see if the value exceeds $2/\sqrt{n}$ and identify the order of the AR model at the lag beyond which partial autocorrelations cease to be significant. The cut-off criterion, where n is the number of shots in the film, is based on the idea that (given appropropriate theoretical conditions) observed $\alpha_h > 2/\sqrt{n}$ differ from zero at the 5% level of significance.

One other idea needs to be explained before examining the methodology of the two papers in detail, and that is *detrending*. Methods of the kind described above for model identification can be invalidated if there are strong trends in the data. One way to avoid problems posed by this is to detrend the data by first fitting an appropriate model of trend and subtracting the values it predicts for the series before proceeding further. Unlike CDN, Redfern (2012b) does this and his approach is emulated in our analyses.

In cinemetric studies prediction is not of interest; rather, the hope is that the order of the AR models identified will provide an insight into the structure of temporal variation of SLs *within* films, and the way structure has changed over a period of time. Specifically, the order of an AR process has been equated with the extent to which shots tend to be clustered with others of similar length, with higher orders associated with greater clustering[1].

Apart from detrending, the other difference in the CDN and Redfern papers lies in the way correlations are measured. CDN measure correlations using the (Pearson's product-moment) correlation coefficient whereas Redfern uses Spearman's rank-order correlation coefficient. This is equivalent to transforming the (detrended) SL data to ranks and is undertaken in the interests of 'robustness', to guard against problems caused by the skewed nature of SL distributions and the presence of outliers. It is argued that will lead to under-estimation of the 'true' order of an AR model. It may be remarked that log-transformation of the data is likely to achieve the same ends; the different possible data treatments do not affect the conclusipns of this paper. Our analysis emulates Redfern in using only 134 of the 150 films studied in CDN, to avoid problems posed by SLs recorded as zero or negative.

Using the `pacf` function in `R`, Figures 1 and 2 show some results for *The Informer* (1935) and *The Grapes of Wrath* (1940). The plot to the left of Figure 1 is for the unmodified SLs, similar to those that form the basis of analysis in CDN. The plot to the right, using rank-transformed data after linear detrending, illustrates Redfern's (2012b) approach. Figure 2 is similar, except that the left-hand plot shows results after detrending the SLs and before rank-transformation.

The horizontal dashed lines in all plots are at $2/\sqrt{n}$, so for *The Informer* would identify AR models of order 0 (to the left) and 2, following the methodologies just described. For *The Grapes of Wrath* orders of 1 and 3 are indicated. Suppressing any concern about these interpretations for a moment, it is sufficient to note that the use of detrending and rank-transformation does make a difference to the results.

CDN express dissatisfaction with this approach for several reasons and develop what they call a modified AR index (mAR). One motivation for this was a feeling that films with smaller values of n, and therefore larger $2/\sqrt{n}$, were 'penalized' in the sense that this 'generated smaller AR indices'. Another motivation, of greater consequence, was based on the observation that 'there can be much noise in partial-autocorrelation functions'. This arguably suggests that for many films an AR model may be inappropriate; however, CDN deal with this issue by using an 'exponential' model to smooth the PACF. This was coupled with a fixed rather than variable bound used to determine, via its intersection with the fitted curve, the mAR index. This approach is adopted

[1] For a discussion of the issue of clustering and reasons for it see the debate between Salt and Cutting on the *Cinemetrics* website, previously referenced, particularly Salt's contribution.

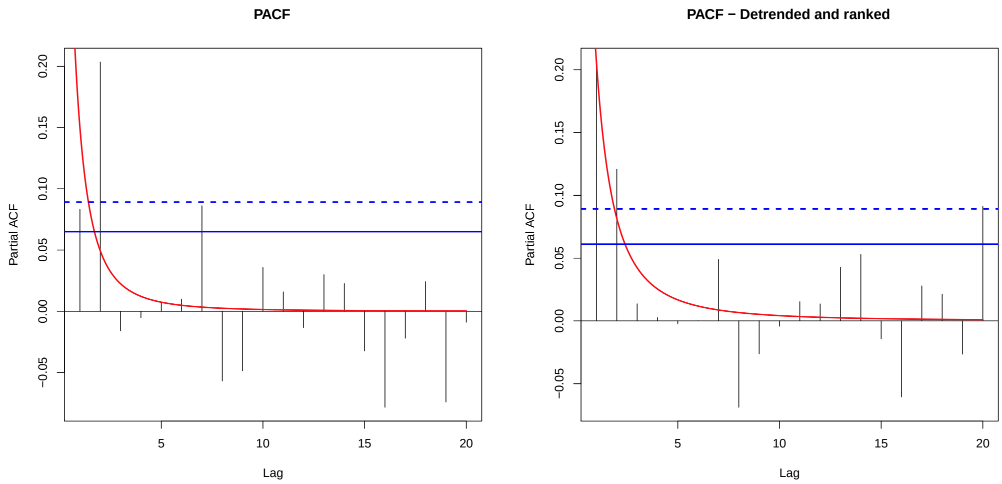

Fig. 1: *PACFs for* The Informer *before linearly detrending the SL data and after linear detrending and a subsequent rank-transformation.*

in Redfern's analysis[2]. This replaces the discrete AR index with a continuous mAR index and introduces the fundamental problem with which much of this paper is concerned.

The fitted lines in both papers, and shown on the graphs, are obtained from an 'exponential' model of the form

$$\alpha_h = (1 + h)^{-\beta}.$$

The `nls` function in `R` was used to fit the models in this paper, with results identical to those in Redfern (2012b).

4 Commentary

This section addresses some purely statistical concerns about the merits of the mAR index. Interpretational issues are considered in the next section.

1. The appropriateness of an AR model is questionable for many of the films. The PACF often bears little resemblance to what might be expected from an 'ideal' AR model (e.g., the left-hand panels of Figures 1 and 2). The ranked data, for the examples shown, are better and this is more generally, though by no means universally, the case.

2. CDN's comment about 'much noise in partial-autocorrelation functions', appears to acknowledged this. Estimating an mAR via the fitting of a smoothed 'exponential' curve disguises rather than deals with the problem of inappropriateness in the first instance. An exponential model is often not appropriate, given the patterns in the data in many cases.

   This can be approached both empirically and theoretically. Most simply, the exponential model is suited to data that, allowing for random variation, exhibit a pattern of decay. None of the plots in the figures really do this. Of the three PACFs shown for illustration in Figure 1 of CDN, fitting an exponential to *King Kong* looks reasonable, but not so *Detour*. If an AR model really is appropriate, then a PACF pattern with a fairly obvious 'drop' is expected, and in this case an exponential model should not be expected to be appropriate. The PACF for *Ordinary People* in Figure 1 of CDN is a case in point; it might be fairly interpreted as

[2] In CDN the mean of n in their sample of films replaces n in $2/\sqrt{n}$; Redfern prefers the median to the mean. These fixed cut-off criteria are represented by the solid horizontal lines in the figures. There is little difference, 0.065 using the mean and 0.061 using the median.

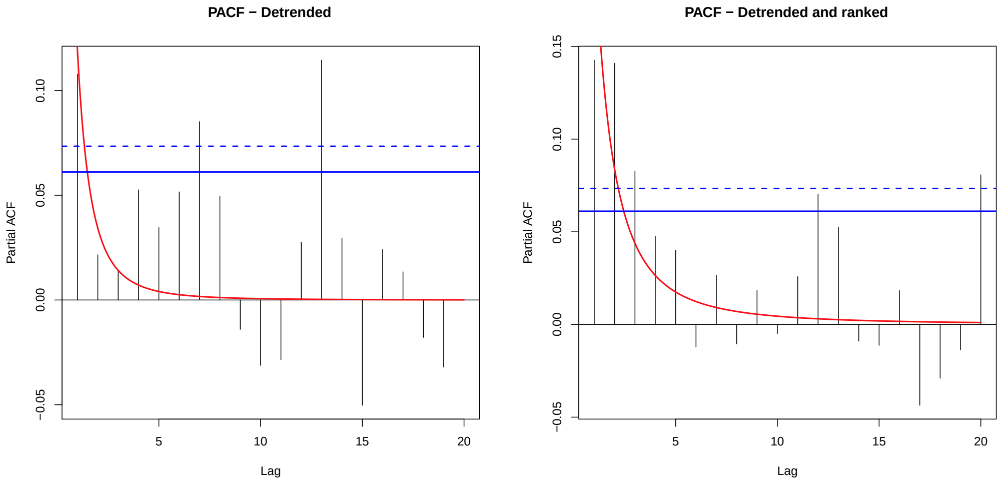

Fig. 2: *PACFs for* The Grapes of Wrath *after linearly detrending the SL data and after a subsequent rank-transformation. Horizontal lines are bounds, used to identify alternative AR indices from a plot.*

consistent with an AR index of 2, but it can then be argued that fitting an inappropriate exponential model to get an mAR of 2.47 is unnecessary.

3. A perhaps minor point is that mAR estimates will be affected by the sample of films used. The mean average shot length in films has been declining over the last 40 years or more, with the consequence that modern films will tend to have rather more shots than older ones from, say, the 1930s and 1940s. This means that while the PACF function for a film will remain unchanged the mAR calculation will depend on the sample of which it is a part, being somewhat larger for a sample of mainly modern films than one spread evenly over a much longer period.

What's being claimed here is that many film SL patterns are not adequately described by an AR model. If they are not, it is difficult to see what the mAR is measuring. If patterns can be adequately described by an AR model the mAR would seem to be redundant. The problem with mARs is, however, worse than this..

5 What does the mAR really measure?

The implication of the above is that interpretations offered for indices for any particular film, and patterns in them over time, can be called into question. The analysis now described was intended to exploit the information in the α_h without imposing on them an interpretation in terms of an underlying AR model, or manipulating them in several stages to get an index of uncertain interpretability. The approach, based on principal component analysis (PCA), led to a rather simple conclusion that allows the PCA methodology that revealed it to be jettisoned[3]. The idea was to apply PCA to the unstandardised α_h, and see if there was any pattern in plots of the PCA scores, against each other, or external factors such as date or genre of the films. This is in the same spirit in which the various AR indices have been used.

[3] PCA is a standard technique described in any good text on multivariate analysis. It produces new variables – principal components or PCs, that are linear combinations of the α_h, that are uncorrelated and explain successively decreasing amounts of variation in the data. It is hoped that some of the PCs are 'interpretable'. Only the first PC was, and it was largely determined by α_1 so patterns can be explored using the latter, forgetting about the mechanism used to reach this conclusion.

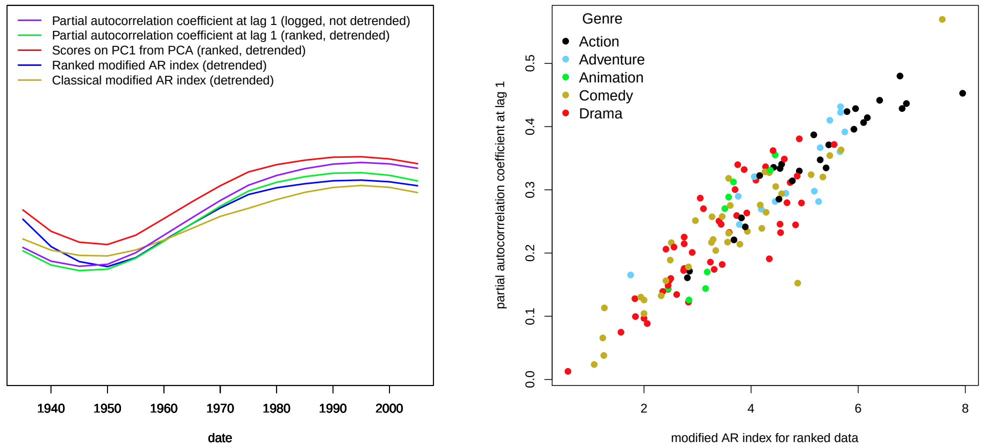

Fig. 3: *To the left plots against date of different 'indices' intended to measure aspects of SL structure in films, for different data treatments and methodologies. Plots are based on loess smooths with a span of 3/4 using localized robust quadratic regression fitting. The right-hand plot is of α_1 and mAR for ranked detrended SLs.*

Only the first principal component (PC1) seemed obviously interpretable; in particular a plot of scores on PC1 against date showed a pattern remarkably similar to that of Figure 3(b) in Redfern (2012b) for his mAR index based on ranked SLs. The pattern doesn't depend on the number of lags used in the PCA. The first PC is largely determined by the value of α_1 and this is all that's needed to obtain the same pattern as in Redfern's analysis.

This is illustrated in the left-hand panel of Figure 3 which contrasts the pattern for the original PCA analysis, the first partial autocorrelation coefficient (of ranked, linearly detrended data), Redfern's (2012b) mAR index for ranked SLs and, for good measure, his mAR index for the detrended but unranked data. Scales for the ordinates are not, for the most part, comparable so have been omitted; plots are overlaid to emphasise the similarity of pattern, but nothing is to be read into the distance between lines[4]. The results for logged SLs, subsequently detrended, are almost identical to those for ranked detrended data and not shown. The curve shown for α_1 for logged data does not involve trend removal of any kind.

An obvious question to ask is why are the results so similar. It suffices to concentrate on a comparison between the first partial autocorrelation coefficient for the ranked detrended data, and the mAR index for the ranked data. The strength of the relationship is illustrated in the right-hand plot of Figure 3, where the genre of a film is also indicated.

Superficially the methodology used to arrive at these very similar results looks rather different. Use of α_1 for showing changing structure over time is pehaps an obvious thought anyway, but was indicated here by the results of the PCA. The PCs are linear combinations of the α_h only the first of which is important here. From the exponetial model previously given, and for a fixed cut-off point, c,

$$\text{mAR} = c^{\beta} - 1$$

so that any variation in the mAR depends on variation in β which is estimated by non-linear least squares methods and is a function of all the 20 lags used. The equivalence between the results suggests that this estimate is very strongly dominated by the magnitude of α_1 but this is less transparent than for the PCA.

[4] The lines are loess smooths using a span of 0.75, and robust quadratic regression models for the localized fitting involved.

Very heuristically, it can be suggested that the larger α_1 is to begin with, the 'longer' the fitted curve is likely to take to fall to its limit of zero, and the larger the mAR, where the curve intersects the fixed cut-off boundary, will be. That this happens can be confirmed via simulation using the `arima.sim` function in `R`, simulating 1000 observations from an AR model of order 1 for different values of α_1, several times for each value. Thus for $\alpha_1 = 0.5$ and ten simulations mAR indices in the range about 3.45-4.00 are obtained – call a typical value 3.7. Reduce α_1 to 0.1 in steps of 0.1 and typical values are 3.1, 2.5, 1.7 and 1.2, with 0.8 for $\alpha_1 = 0.05$.

This is a very small experiment and there is a lot of variation about the typical values quoted, but the pattern is clear enough. If 134 such analyses were to be performed, varying α_1 between 0.05 and 0.50, and assigning 'dates' to each simulation with later dates corresponding to higher α_1, plots like those to the left of Figure 3 and more linear might be expected. This would show an 'evolution' in mARs that, if the interpretation of CDN was accepted, would be interpreted as an increase in clustering into packets of shots of similar length.

It cannot be so interpreted. In CDN the AR index is replaced by the mAR to deal with the fact that the real data often does not look like an AR process. By contrast, here the simulated data, by construction, are from an AR process, so there is no need for an mAR and no 'evolution' since the processes are all of the same order. The suggestion of 'evolution' in the simulated example is entirely an artefact of the methodology used to construct mARs, whereby larger α_1 generate larger mARs, regardless of the true order of the underlying process.

The above is based on the simulation of AR models of order 1. It is more tedious to investigate models of higher order, as a lot more variation in the parameters that define a model is possible. I have not looked at this systematically, but it is easy enough, using simulated data, to show that mARs more nearly reflect the values of the larger partial autocorrelations present than the order of the underlying process. As a rough illustration of the kind of variation to expect, using detrended and ranked data, there are 49 films with an AR index of 2, whose mARs range from 1.94 to 7.57, and 21 films with an AR index of 3 whose mARs vary between 2.45 and 5.79.

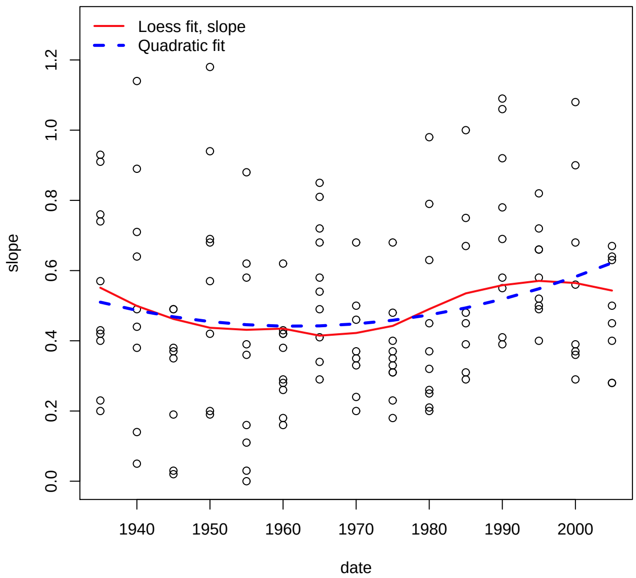

Fig. 4: *Plots of α ('slope') against date using the model of $1/f$ patterning developed in Cutting* et al. *(2010) and a loess smooth of the same data. See the text for detail.*

Only a brief discussion of $1/f$ analyses is attempted, using the results reported in the supplement to CDN. They smooth the power spectrum for the (standardized) SLs of a film using 7-10 estimated points, to which a line determined by a parameter, α, is fitted. The closer α is to 1 the closer the film is to exhibiting $1/f$ patterning. The estimated values of α are plotted against date, a line through the points fitted with a quadratic regression model, and 'evolution' towards $1/f$ patterning claimed on the basis that the right 'arm' of the fitted model is moving smoothly upward towards 1. This is illustrated in Figure 4 along with a loess smooth to the same data, of the kind used in the previous figure[5].

It can be remarked in passing that, on the basis of Figure 3 in CDN, the propriety of fitting a linear model to estimate α for some films can be questioned; the data do not have the shape to merit such a fit. Even so, and ignoring this, the loess smooth to the data in Figure 4 shows no sign of 'evolution', the short period beween about 1970 and 1990 being the only one when anything of the sort was happening. The question has been asked (by Salt) as to why, if there has been an evolution to $1/f$ patterning since 1935, we aren't closer to it. A perfectly plausible answer is that no such evolution is taking place. I am not claiming that the loess smooth is 'right' and the quadratic smooth 'wrong' – there is too much variation in the data for either to be regarded as compelling evidence for a 'trend' of any kind. That the data can tolerate this kind of variation in the fitted lines, though, is reasonably compelling evidence that claims for an evolution to $1/f$ patterning cannot be strongly pressed, if at all.

6 Discussion

The scope and ambition of the work initiated in CDN is impressive, technically demanding at the statistical level, and leads to some 'eye-catching' conclusions. Their paper concluded with the claim that the study had 'demonstrated that the shot structure in films has been evolving toward $1/f$ spectra', liberally interpreted in parts of the media as relating to the secrets of making a 'Hollywood blockbuster'. More seriously, and soberly, the claims of the CDN paper have been repeated in subsequent academic literature. The present paper shows that they cannot be sustained.

As far as the evidence for 'evolution' in terms of clustering into packets of shots of similar length goes, the underlying assumption that fitting AR models is appropriate seems not to have been questioned. The comment in CDN that PACFs can exhibit a lot of 'noise' might be read as a tacit admission that there are problems. Even if the mAR can be interpreted as a continuous analog of the AR index it is of no value if the process is not AR to begin with.

More seriously, even if an AR model is valid, the mAR does not measure what is claimed for it. Rather than measuring the order of an AR process, interpreted in turn as a measure of the degree of clustering, what it actually appears to be is a surrogate measure for the (partial) autocorrelation at the first lag. This only provides information on whether the process is of order zero or not. Processes of the same order can give rise to entirely different mARs, and processes of different orders can give rise to the same mAR. The mAR is of no value as a diagnostic tool for the order of an AR process.

What this means is that the interpretation of analyses based on the patterning of mARs over time - whatever these patterns might be - that treats the mAR as indicative of clustering in the manner claimed cannot be sustained. Neither the mAR pattern (however the mAR is interpreted) nor that of the slopes associated with $1/f$ ratios, can be unequivocally interpreted as indicative of evolution if by this is meant a continuing upward trend in the indices involved. If the data are allowed to speak without having a story (i.e. linear or quadratic fit) imposed on them their tale, if any, is of evolution having 'stalled' over the last 20 or 30 years.

Figures 3 and 4 do exhibit pattern but not of a straightforward evolutionary kind. The patterns in Figure 3 can be interpreted in terms of variation in the first-order autocorrelation coefficient. The interpretation of this raised its head in interchanges in the Salt/Cutting debate. The patterns over time for the mAR and $1/f$ statistics *claimed* in CDN mirror, inversely, those of the ASL which has been 'evolving' in the form of a fairly steady decrease, in its mean, over most of the period

[5] For consistency the 134 films analyzed elsewhere in this paper have been used. There are slight differences between the figure and Figure 2(c) in CDN as a consequence, but if all 150 are used the loess curve suggests, even more strongly, that 'evolution' is not occurring.

from 1950 on. Salt wondered about the nature of the relationship between the ASL, mAR and $1/f$ statistics, relating this to variation in the Lag-1 autocorrelation. Cutting, in his reply, stated that 'Lag-1 data alone can be misleading' but his subsequent advocacy of the mAR as a superior measure of clustering fails with the demonstration that the mAR is essentially a surrogate for the Lag-1 autocorrelation. Salt proposed an heuristic argument, centered on action films, that there were 'likely to be more strings of shots of nearly the same length' (i.e. clustering) as the ASL decreased. The whole discussion, though, is predicated on the assumption that evolution of some kind is taking place, and while this may be true for the ASL the argument presented in this paper is that it is not true of the mAR (Lag-1 autocorrelation) or $1/f$ ratio.

The pattern of mean ASL against year for US feature films is shown in Salt (2009, p. 278). The Lag-1 autocorrelation might be re-interpreted as an index of clustering, and in a sense this would 'rescue' one of the interpretive tools in CDN by replacing the mAR with the Lag-1 autocorrelation so interpreted. However, Salt's argument for an association between clustering and the ASL only 'works' up to about 1990 and for mean ASLs down to about 6.5. The mean ASLs vary around 7 to 6.5, with a slight decline from about the mid-1970s to 1990 then show a sharp decrease. This sharp decrease is not mirrored by an increase in the patterns in Figures 3 and 4. Thus, conclusions about 'evolution', of the form that have been claimed, are unsustainable even if we allow that the mAR is a measure of clustering (albeit not in the manner intended in CDN). In the absence of evolutionary patterns, attempts at rationalizing their existence in terms of theories of attention need to be put on hold. The more interesting problem might be to explain why evolution that mirrors that of the ASL is not taking place.

References

Abel, R. 1989: Split Decision. *Quarterly Review of Film and Video* **11**, 43-57.

Bosse, A., Tsivian, Y., and Brisson, K. 2011: Cinemetrics: A digital laboratory for film studies. http://dh2011abstracts.stanford.edu/xtf/view?docId=tei/ab-368.xml;query=Arno%20Bosse;brand=default (accessed 5 November 2013).

Cutting, J.E., DeLong, J.E. and Nothelfer, C.E. 2010: Attention and the evolution of Hollywood film. *Psychological Science***21**, 440-447.

Redfern, N. 2012a: The average shot length as a statistic of film style. http://www.cinemetrics.lv/dev/on_statistics.php (accessed 5 November 2013).

Redfern, N. 2012: Robust estimation of the modified autoregressive index for high grossing films at the US box office, 1935 to 2005. http://nickredfern.files.wordpress.com/2012/11/nick-redfern-the-mar-index-for-hollywood-films1.pdf (accessed 5 November 2013).

Salt, B. 1974: Statistical style analysis of motion pictures. *Film Quarterly* **28**, 13-22.

Salt, B. 2006: *Moving Into Pictures.* London: Starword.

Salt, B. 2009: *Film Style & Technology: History & Analysis (3rd edition).* London: Starword.

Smith, T.J., Levin, D. and Cutting, J.E. 2012: A window on reality: perceiving edited reality. *Current Directions in Psychological Science* **21**, 107-113.