More on the Evolution of Popular Film Editing

James E. Cutting
Department of Psychology, 109 Tower Road, Cornell University
Ithaca, NY 14853-7601 USA
(jamescutting@cornell.edu)

My students and I have been pleased to see that the results of Cutting, DeLong, and Nothelfer (2010) have received considerable attention – in the legitimate press, on blogs (although often mis-citing our conclusions), in academic journals, and within in the cinemetrics community. With respect to the latter, Barry Salt (2010a, 2010b) and I (Cutting, 2010a, 2010b) have already had exchanges on this, and now Mike Baxter (2013) has joined in, along with some new thoughts by Salt (2014).

Autoregression. Let me first address what I take to be the major issue raised by Baxter, who focused on our invention and use of a modified autoregressive model (mAR). Autoregression is a procedure, used mostly in signal processing, to look at the patterns in a signal (a vector, or series of numbers) by comparing it to itself at various lags. A lag of 0 indicates a straightforward correlation of all values with themselves, and the correlation is always 1.0. However, things begin to get interesting when one staggers the values by 1, correlating the 1st with the 2nd value, the 2nd with the 3rd, and so forth. A significant correlation at Lag 1 indicates a carryover effect and a nonindependence in the series of numbers. And, as one might suspect, one doesn't stop there, but goes on autocorrelating at Lags 2, 3, 4, and so forth. Of interest here, is the sustained significant correlations as the lags increase. The longer the stretch of significant correlations the more the elements of the signal (here shot durations) effect one another, indicating what we have called the systematic and increasing "local" effects in the pacing in movies.

How does one assess? One first asks if the autocorrelations are reliable at Lag 1; if so are they reliable at Lag 2 and still reliable at Lag 1? If so are they reliable at Lag 3 and still reliable at Lags 1 and 2? And so forth. One continues this process until one has reached the maximum lag that is reliable with all the previous lags also reliable. Typically in signal processing one detrends the data, and this is what Baxter (2013) and Redfern (2012) have done. Detrending is done, normally, to make sure that the ends of the vector (the first and last numbers) are the roughly same, extracting a regression line from the data, and linearly changing all of the data in between. It should be said that movies have a typical, and different trend that is far from linear (Cutting, Brunick, & DeLong, 2012) – the shots of the first few minutes of a film tend to be longer than all others, followed by a long plateau that goes through the ¾ mark of the film (to the beginning of the climax); the shots then grow shorter still, through the climax, and then return to the general level of the film average during an epilog.

The autoregressive procedure essentially treats the numbers as a circular vector, abutting the first and the last. If there is a great discontinuity between the ends, and particularly if the vector length is relatively short, the results without detrending can be very misleading. However, if the vector is long and relatively stationary (the variance doesn't change much throughout) detrending is not really necessary. Moreover, when I detrended a selection of these films, the results I am about to report didn't change. Thus, we didn't detrend these data in Cutting et al (2010), and I haven't detrended them in the new analyses here.

Back to mAR. We invented mAR in part, as Baxter noted and as we said in Cutting et al (2010), because partial autocorrelation functions are noisy (see also Redfern, 2012). We also wanted a measure that was continuous rather than discrete because the statistics we would use assumed continuous variables. Nonetheless, the careful analyses of Baxter suggest that the decision to use mAR was not entirely effective in capturing the essence of the partial autocorrelation functions. More pertinently here, however, the left panel of his Figure 3 shows that, for various measures of autocorrelation, detrended or not, the pattern of results is essentially the same. For our film results there is a curvy line descending from 1935 to about 1960 then arching upward with an apparent asymptote at about 1990. The confluence of Baxter's five separate methods is a testament to the robustness of the data and the noncriticality of any particular method. Moreover, looking at his figure, it seemed to me that it was reasonably plausible that there was still a linear trend lurking there.

So, using the statistical package JMP I reanalyzed the data of the 150 films reported in Cutting et al (2010) and the ten newer films released in 2010 reported in Cutting, DeLong, Brunick, Iricinschi, and Candan (2011). As did Salt (2014), we corrected 14 of the shot-duration series in the first group (not analyzed by Redfern, 2012 or Baxter, 2013) to remove the negative and zero shot durations (due to typos in the column of frame numbers). Moreover, some of these series are different than those that we deposited on the Cinemetrics website. As we go through these films for new analyses, we occasionally discover cuts that we missed or that we falsely inserted before. Thus, in my new analyses I include the corrections made for the 24 films analyzed by Cutting, Brunick, and Candan (2012) and Cutting and Iricinschi (2014). The admonitions of Smith and Henderson (2008) in their discussion of "edit blindness" (not detecting cuts) should be taken seriously by cinemetricians – it is a very difficult task to determine all the shots of a complete film.

For their AR analyses Baxter and Redfern used ranked data. That is, the magnitudes of each shot duration were ranked and then the analysis done on the ranks rather than on the normalized (mean = 0, standard deviation = 1) metric data. A reason for this is to remove the effects of significant outliers (long takes). But Baxter also suggested that a log transform of the data might satisfactorily remove such outliers, so I used the log-scaled shot-duration data for each film and then normalized them. The procedure was as outlined above. I incrementally increased the integer in the specified AR(*n*) algorithm until I lost a string of significant results.

I then recorded the integer *n* that yielded a significant correlation with every lower lag (*n* - 1, *n* – 2, etc.) also significant.

Data are shown in Figure 1, with the AR index (the largest significant string of correlations) on the ordinate (y axis) and release year of the film on the abscissa (x axis). Remember, the results of the AR(*n*) analysis only allow for cardinal results for *n* – 0, 1, 2, 3, 4, 5, 6, and even 7. Thus, clusters of films for each release year and each AR value occupy the same space in the graph, but are spread here to see the groupings of identical results for different movies. In what follows I will assume that these are continuous data. Obviously they are not, but I cannot make the comparisons needed without this assumption. It is also an assumption that is made quite a bit in psychology, and particularly psychophysics, since analyses on large sets of ordinal data often converge strongly on continuous (metric) data.

A linear model fits the data well ($t(158) = 6.85$, $p < .0001$, $d = 1.09$), accounting for 22.89% of the variance in the data. The cubic model fits somewhat better ($\Delta F(2,158) = 3.21$, $p = .043$, $d = .28$), accounting for 24.99% of the data. The cubic model is somewhat like Baxter's loess (or lowess, for LOcally WEighted

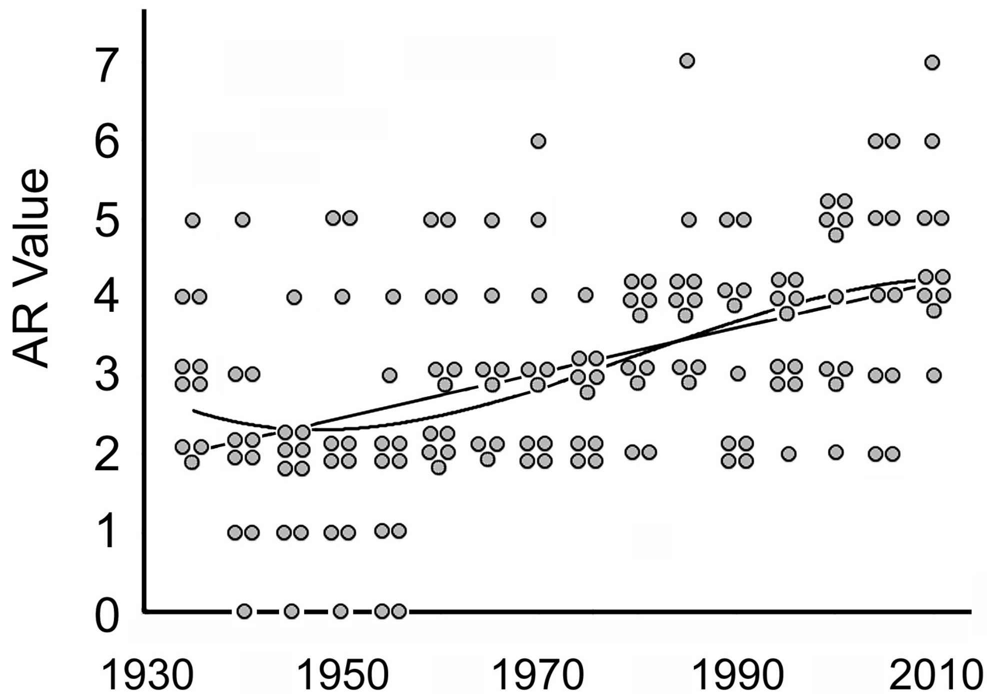

Figure 1: A scatterplot of the autoregressive indices – AR(n) – for 160 popular films, 10 in each of 16 release years from 1935 to 2010 in five year intervals. Two fits to the data are also shown: a linear fit (endorsed originally by Cutting, DeLong, and Nothelfer, 2010) and a cubic fit, with roughly the same shape as the loess fit used by Baxter (2013).

regreSSion) fit with a quadratic kernel and moving window of ¾ of the total release-year span. The reasons for choosing the cubic fit here are twofold – first, it seems to mimic well Baxter's loess fit in his Figure 3 (left panel); and second, it allows a comparison using the number of parameters in the two models. Loess fits are descriptive, but because it can be difficult to define their parameters they are less appropriate for model comparisons.[1]

Since the models have unequal numbers of parameters (linear = 1; cubic = 3) it is traditional in the model fitting literature to penalize the model with more parameters (see, for example, Myung & Pitt, 2002). The reason is that the more complex model might be overfitting the data. The rationale for this is fairly straightforward; two points can be fit by a line (n-1, or 2-1, or 1 dimension), three points by a plane (3-1 or 2 dimensions), and so forth. Generalizing, any model with n-1 parameters (dimensions) can fit n data perfectly. This is overfitting. Of course, most models have nowhere near the number of data being fit, but this is always an important issue to consider.

Comparisons among models with different numbers of parameters can be done by considering the Bayesian Information Criterion (BIC) for each fit. This is a calculation where the model with the smaller BIC number is considered the better fit in light of parameter differences. The BIC value for the linear fit is 531.42, whereas that for the cubic fit is 537.15. The magnitude of these numbers doesn't matter, only the fact that one is smaller than the other. Thus, the linear fit should be considered better that the cubic fit. Given the closeness of the comparison seen in Figure 1, however, only a curmudgeon would soundly reject the cubic fit. Nonetheless, at least there is certainly no reason to reject the linear fit in this context.

Salt (2013) is correct about the relations among AR values, release years, and number of shots per film. That is, when both shots and release year are used in a regression to predict the AR index, the effect of shots is reliable ($t(157) = 9.54$, $p < .0001$, $d = 1.52$) whereas the effect of release year is not ($t(157) = .89$, $p = .37$). Thus, insofar as I can claim that there is an evolution in the increasing correlations in the shot durations of popular movies, this seems driven by the fact of shorter shot durations, and thus more shots. The importance of the both AR index trends is that it measures the correlations of shot durations "locally" in the vector of shot durations.

Fractal-like patterns in the shot durations of movies. Cutting et al (2010) analyzed the shot-duration series in two ways – "locally" (using autoregressive measures) and "globally" (using a whitened fractional Brownian measure, a combination of $1/f^\alpha$ and white noise; Gilden, 2005, 2009). We did so because there is considerably controversy over local vs. global effects in the literature on

[1] Baxter's quadratic kernel has two parameters and, intuitively, the ¾ moving window along the shot vector would be a third.

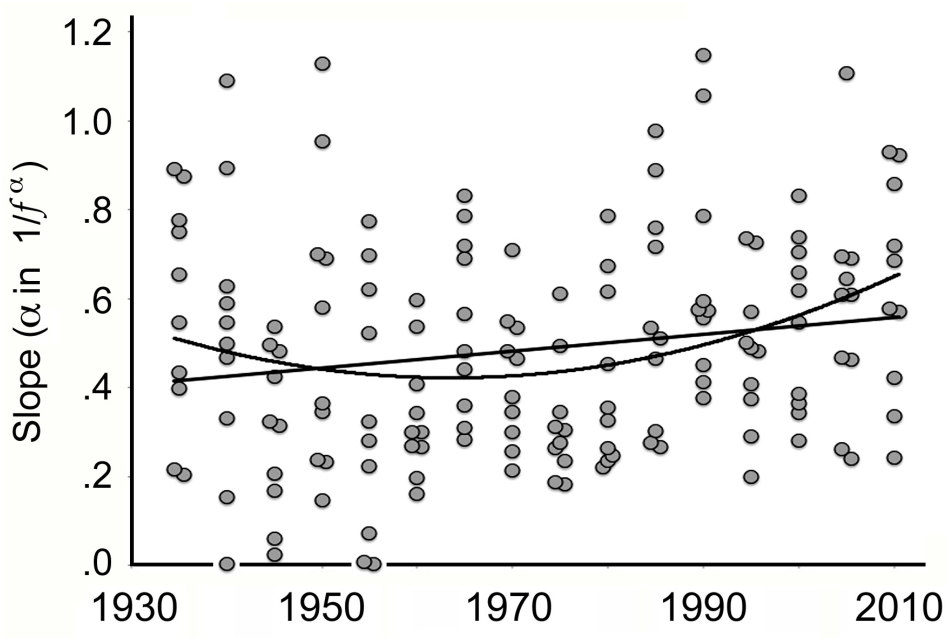

Figure 2: A scatterplot of the slope (α) of each of the same 160 popular films where each series of shot durations was fit to a function with two parameters, white noise and a $1/f$^a component (Gilden, 2001; Thornton & Gilden, 2005). The scatterplot is fit with both linear function and quadratic functions.

serial reaction time (RT; Van Orden, Holden, & Turvey, 2005; Wagenmakers, Farrell, & Ratcliff, 2004). We hoped that by using both, and by finding similar trends in both, we could set this controversy aside.

Baxter (2013) spent little space reanalyzing the differential slopes of popular films reported by Cutting et al (2010), which we regarded as the most important result in our paper, and which received the most press. We claimed that these slopes – the exponent α the pattern $1/f$^α fit to the power spectrum derived from the series of shot durations – increasingly approached a value of 1.0 from about 1960 to the present day. Both linear and quadratic trends were reliable. Moreover, what made that result potentially newsworthy is that the pattern of reaction times in many cognitive tasks also has an exponent α that approaches 1.0. Thus, perhaps movies, which exogenously drive our attention through eye movements that occur after every cut (Smith, Levin, & Cutting, 2012) are evolving to match the natural, endogenous patterns of long temporal spans of our attention.

Baxter refit our slope data with another loess function in his Figure 4, and found, at least to his eye, less than convincing evidence for an upturn in a quadratic function. Fair enough. So I recalculated the slopes for all of the films, especially including those with slightly modified shot duration vectors, and also added the ten films from 2010. Results are shown in Figure 2, with linear and quadratic regression lines fit to the data.

The linear function, although very shallow, is statistically reliable ($F(1,159) = 9.64$, $p < .024$, $d = .49$). Thus, although it would likely convince only the faithful, there is an increase in slope over release years, suggesting an evolution in the editing patterns of popular movies. Moreover, this increase remains reliable ($t(157) = 2.06$, $p < .04$, $d = .33$) when regressed with the number of shots per film, which is also a reliable predictor. Notice, again in congruence with the suspicions of Salt (2014), the effect of the number shots is the stronger of the two ($t(157) = 5.24$, $p < .0001$, $d = .84$).

The linear function accounts for 3.17% of the variance in the data, whereas the quadratic function accounts for 7.32%. The BIC criterion favors the quadratic fit (BIC = 12.80) over the linear fit (14.73). Thus, the quadratic fit to the data should be preferred despite its extra parameter (2 vs. 1). Both results parallel those found in Cutting et al (2010). Thus, I see no reason to step back from our previous conclusion: Our study "demonstrated that the shot structure in film has been evolving toward $1/f$ spectra (again, mixed with white noise)" (Cutting et al, 2010, p. 7). Again, the importance of the slope measure is to demonstrate the change in the "global" pattern of correlations.

Salt (2014) doesn't much like our two-component fits of white (β, or random) noise and $1/f\,^\wedge\alpha$ (non-random) noise spectra. He asks (p. 2): "how does one identify which values relate to which part of the division into random and non-random values." The answer is that one simultaneously fits both α and β to the data, and finds the best simultaneous fit. This is a fairly standard procedure in model fitting. But more importantly, he worries (p. 2) that "all shot lengths are signal, and none of them are noise, since their [sic] was a series of more or less conscious choices to make the shots the length they are." Of course, there is a sense in which Salt is correct, but there is an equal sense in which RTs are also all signal. The issue is that the terms "signal" (or $1/f\,^\wedge\alpha$ pattern) and "noise" are used in special ways in this context. In fact, both are called "noise" with the former called "pink" and the latter called "white."

Gilden's (2001, 2009) has called his model *whitened fractional Brownian motion.* The "whitened" part is the addition of white noise, and the "fractional Brownian" part means that the term α in $1/f\,^\wedge\alpha$ can vary (but will typically be near, and often below, 1.0). In its application to RTs, these two components have well established precedents. Wing and Kristofferson (1973) modeled RTs with two components, one of neural noise (what state the nervous system is in in its readiness to respond, which can be modeled as random, or as "white noise") and one of motor output (for example, how long it takes to initiate and execute a finger-press, non-random). Gilden (2001, 2009) simply substituted $1/f\,^\wedge\alpha$ noise for the nonrandom part.

What are these random and non-random components in a series of shot lengths? Good question. One pragmatic answer is that these needn't be specified since Cutting et al (2010) were simply looking for a parallel between the series of shot durations and those of RTs. A more acceptable answer, however, could be that the random component can be found in the juxtaposition of details within the takes that must appear before a cut can be made. Suppose one is choosing a shot among ten takes for a given shot. The critical portions of all of the takes will be different in duration and the differences can be considered more or less random (β) around some mean value. The "ideal" of the critical attributes of the take, which would form the basis of why the shot would appear in the movie, could be the mean duration of all of those takes. This value could contribute to global component (which happens to be modeled by $1/f^\alpha$).

In other words, every shot has to do something in a movie, otherwise it wouldn't be there. That something will vary in duration across takes, where β represents that variance (with the assumption that that variation is random) around some mean. The editor selects the frames from one of these takes which will have a duration of the average sequence in the takes, plus or minus some random interval. In other words, one might say that the $1/f^\alpha$ component measures the conceptual intentions of the film editor (and director) along with the narrative structure of the movie, whereas the white noise component measures the variability of conditions across the set of takes. And this is not the only possible construal of a two component editing process, but it seems reasonable.

Scientific Inference. In summary, it might prove useful to reconsider Baxter's (2013) four-step approach to our analysis (Cutting et al, 2010), with my annotations below:

> 1. The shot duration data for a film are converted, by statistical methods, into indices measuring some property of interest.

Cutting et al (2010) did this and I've done it again here, this time for AR indices (a "local" measure in the correlations of shot durations), and again for slope indices (α, a "global" measure of correlations). In both analyses some of the data are corrected from those used by Cutting et al (2010); again, it would appear that because of "edit blindness" (Smith & Henderson, 2008) one must be continually wary of statements that complete accuracy in determining shot structure has been attained.

> 2. Statistical methods are used to identify temporal patterning in the indices.

Again, I have repeated this, and the new analyses reveal upward trending patterns for both indices, from at least the 1950s to the present for the AR measures (a local measurement) and at least the 1970s to the present for the slope measures (a more global measurement). Statistically appropriate claims, however, can be made for linearly increasing trends from 1935 to 2010 for both measures.

3. A 'filmic' interpretation is offered for the patterns detected, namely they are evolving.

Evolution in this context means consistent, directional change; and might be opposed to faddishness, where changes oscillate among alternatives. I have deliberately conflated the use of this one meaning of the term "evolution" with its biological meaning. Clearly people are not evolving to watch movies, but it can be said that movies are changing to accommodate perceptual and cognitive abilities of viewers. It should also be said, however, that that statement was largely speculation in 2010. Nonetheless, the claim seemed to make intuitive sense.

More recently, my students and I have tracked many other "evolutionary" changes in movies in our sample over the span from 1935 to 2010 (Cutting & Candan, 2013), putting more meat on the general claim. Beyond the generally and roughly linear log-scaled shortening of shot durations, certainly the most commonly cited change in film (Bordwell, 2006; Cutting, Brunick, & Candan, 2011; Cutting, DeLong, Brunick, Iricinschi, & Candan, 2001; Salt, 2006, 2009), there has been a more or less linear increase in motion and change in movies (Cutting, DeLong, & Brunick, 2011; Cutting, DeLong, Brunick, Iricinschi, & Candan, 2011), a more or less linear increase in the (negative) correlation between shot durations and the amount of motion in the shot (shorter shots now have more motion; Cutting, DeLong, Brunick, Iricinschi, & Candan, 2011), a generally linear shift shot scales towards more closeups (Salt, 2009; Cutting & Iricinschi, 2014; Cutting, 1914a), an increase in the use of luminance and motion contrasts across the screen (Cutting, 2014b), a change in the use of location, time, and character shifts across scenes (Cutting, 2014a), a linear shortening of shot durations and shot scales in re-establishing shots (scene-opening shots revisiting locations seen before; Cutting & Iricinschi, 2014), and a decline in the use of dissolves and fades (Carey, 1974, 1982; Cutting, Brunick, & DeLong, 2011). It should now be clear, then, that all of these consort to support the idea that film style is evolving.

4. Psychological theories of attention are offered to 'explain' this evolution.

Every cut demands a reallocation of attention (that is, at minimum, it will cause the viewers eye to move). About 95% of all shot transitions in the corpus of films sampled are separated by cuts, and in contemporary films that value is almost 99%.

It happens that endogenous attention patterns, as measured in RTs in perceptual and cognitive experiments, generate RT series that have a $1/f$-like pattern. It also happens that the shot-duration patterns in the films analyzed in this corpus can be said to be approaching a $1/f$ pattern.

These statements reveal only a parallel; they do not really form an "explanation," or at least not causal connection. The linkages are these: (a) Natural attention over an extended timeframe is patterned, (b) attention can be driven by eye movements, (c) cuts in film generate eye movements, and (d) cut patterns in film generate patterns that are increasingly like those of natural attention. Again, there is no logical necessity for accepting this parallel and our conclusion; no causal linkage has yet been demonstrated (although we are working on this). But, I would claim again, it does have some intuitive plausibility.

Cross-disciplinary differences, sampling, and data. Data are data and can be fit in many ways. I have no problem with loess fits or with moving averages as *descriptions* of data. However, I do have problems when they are used in conjunction with models and theories. The problem, at least for me as a psychologist, is that I have a modicum of distrust in our own sampling methods when we chose these 160 movies.

Not trusting ourselves to pick films, my students and I used an external source – IMDb (Internet Movie Database) ratings and reports – to select movies and we chose those from among the most popular and the most often seen, under the constraint that we wanted movies from five genres, if they could be had – animations, action films, adventure films, comedies, and dramas. But did we choose truly “representative” movies of each release year? I have no idea and I will never be sure; moreover I don’t even know what “representative” might mean. The upshot of this quandary is that when I look at the small wrinkles in the loess and moving average fits such as those plotted by Baxter (Figure 4) my first response is to think that the stimulus movies may have been slightly mis-sampled. After all the variance differences in slopes in Figure 2 across different release years is impressive. My choice of using linear, quadratic, and cubic model fits is in the hope of glossing over slightly misrepresentative subsamples in any given release year.

Another attraction of linear, quadratic (and even cubic) fits is that they seem “simple” and cover the entire data set. Simplicity in theories has long been a popular tenet in science (Goodman, 1983; Kuhn, 1977), and is reflected in the increased use of metrics such as BIC in model comparisons. One could argue that this is a belief in some Platonic truth about data and results; that some “ideal” and “truthful” function underlies the mess of reality. But actually, at least for me, it is more of a disbelief in relatively small numbers. Ten films per release year is not very many, but before Cutting et al (2010) no one had analyzed even one film in the manner that we had.

In conclusion and in brief, I am pleased for Baxter’s careful consideration of our data and for Salt’s (2014) continued grappling with the issues we have broached, but after reanalyses of our data I see no reason to change from our original stance.

References

Baxter, M. (Nov 2013). Evolution in Hollywood editing patterns? Cinemetrics Website. http://www.cinemetrics.lv/dev/Evolution_paper_for_Cinemetrics.pdf

Bordwell, D. (2006). *The way Hollywood tells it*. Berkeley: University of California Press.

Carey, J. (1974). Temporal and spatial transitions in American fiction films. *Studies in the Anthropology of Visual Communication, 1*, 45-50.

Carey, J. (1982). Conventions and meaning in film. In S. Thomas (Ed.), *Film/culture: Explorations of cinema in its social context* (pp. 110-125). Metuchen, NJ: Scarecrow Press.

Cutting, J. E. (2014a). Event segmentation and seven types of narrative discontinuity in popular movies. *Acta Psychologica, 149*(6), 69-77.

Cutting, J. E. (2014b). How light and motion bathe the silver screen. *Psychology of Aesthetics, Creativity, and the Arts,* in press.

Cutting, J. E. (2010a). In reply to Barry Salt on *Attention and the Evolution of Hollywood film.* Cinemetrics Website. http://www.cinemetrics.lv/cutting_on_salt.php

Cutting, J. E. (2010b). On Salt II. Cinemetrics Website. http://www.cinemetrics.lv/cutting_on_salt_on_cutting_on_salt.php

Cutting, J. E., Brunick, K. L., & Candan, A. (2012). Perceiving event dynamics and parsing Hollywood films. *Journal of Experimental Psychology: Human Perception and Performance, 38*(6), 1476-1490.

Cutting, J. E., Brunick, K.L. & DeLong, J. E. (2011). The changing poetics of the dissolve in Hollywood film. *Empirical Studies in the Arts, 29*(2), 149-169.

Cutting, J. E., Brunick, K. L., & DeLong. (2012). On shot lengths and film acts: A revised view. *Projections: The journal for movies and mind, 6,* 142-145.

Cutting, J. E. & Candan, A. (2013). Movies, evolution, and mind: Fragmentation and continuity. *The Evolutionary Review, 4*(3), 25-35.

Cutting, J. E., DeLong, J.E., & Brunick, K. L. (2011). Visual activity in Hollywood film: 1935 to 2005 and beyond. *Psychology of Aesthetics, Creativity, and the Arts, 5,* 115-125.

Cutting, J. E., DeLong, J. E, Brunick, K. L., Iricinschi, C., & Candan, A. (2011). Quicker, faster, darker: Changes in Hollywood film over 75 years. *i-Perception, 2*, 569-576.

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

Cutting, J. E. & Iricinschi, C. (2014). Re-Presentations of space in Hollywood films: An event-indexing analysis. *Cognitive Science,* in press.

Gilden, D. L. (2001). Cognitive emissions of 1/*f* noise. *Psychological Review, 108(1)*, 33-56.

Gilden, D. L. (2009). Global analysis of cognitive variability. *Cognitive Science, 33*(8), 1441-1467.

Goodman, N. (1983). *Fact, fiction, and forecast.* Cambridge, MA: Harvard University Press.

Kuhn, T. S. (1977). *The essential tension: Selected studies in scientific tradition and change*. Chicago: University of Chicago Press.

Pitt, M. A., & Myung, I. J. (2002). When a good fit can be bad. *Trends in Cognitive Science, 6*, 421-425.

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

Salt, B. (2006). *Moving into pictures*. London: Starword.

Salt, B. (2009). *Film style and technology: History and analysis* (3rd ed.). London: Starword.

Salt, B. (2010a). Comments on *Attention and Hollywood films.* Cinemetrics Website. http://www.cinemetrics.lv/salt_on_cutting.php
Salt, B. (2010b). Some more thoughts on the Cutting article. Cinemetrics Website. http://www.cinemetrics.lv/salt_on_cutting_on_salt.php
Smith, T. J., & Henderson, J. M. (2008). Edit blindness: The relationship between attention and global change blindness in dynamic scenes. *Journal of Eye Movement Research, 2,* 1–17.
Smith, T. J., Levin, D. T., Cutting, J. E. (2012). A window on reality: Perceiving edited moving images. *Current Directions in Psychological Science, 21*(2), 107-113.
Thornton, T.L. & Gilden, D.L. (2005). Provenance of correlations in psychological data. *Psychonomic Bulletin & Review, 12(3)*, 409-441.
Van Orden, G. C., Holden, J. G., & Turvey, M. T. (2005). Human cognition and 1 / ƒ scaling. *Journal of Experimental Psychology: General, 134*(1), 117–123.
Wagenmakers, E.-J., Farrell, S., & Ratcliff, R. (2004). Estimation and interpretation of 1/*f* ^a noise in human cognition. *Psychonomic Bulletin & Review, 11*(4), 579–615.
Wing, A. M., & Kristofferson, A. B. (1973). Response delays and the timing of discrete motor responses. *Perception & Psychophysics, 14*, 5–12.