A wonderful night for Oscar speeches

Winning an Academy Award is a career highlight for actors, directors and anyone else involved in the movie business. But the happiness and gratitude that accompany a win often lead to long, rambling acceptance speeches. **Adam B. Kashlak** analyses speech lengths and ceremony runtimes since 1942

Adam B. Kashlak is an assistant professor of statistics at the University of Alberta. He completed a PhD at the University of Cambridge and was a statistician at the US Department of Defense.

Sunday, 24 February 2019 will be – as actor Billy Crystal often sang – "a wonderful night for Oscar". The glitz and glamour of Hollywood will descend upon the Dolby Theatre as the 91st Academy Awards ceremony begins. With it comes an evening of song, dance, jokes, laughter – and speeches. Very long speeches.

Oscar acceptance speeches have reached a level of infamy for being rambling emotional expulsions, often served with a dash of political ire or condescending praise for the so-called "little people" (to quote Barbra Streisand). The producers of the awards try to rein in the winners by cutting mics and cuing the music, as each speech is supposed to be limited

to 45 seconds. Last year, host Jimmy Kimmel even offered an $18 000 jet ski to the person giving the shortest speech, though his efforts appear to have been in vain.

The Academy Awards ceremony has certainly expanded from its 1927/28 debut, which lasted a mere 15 minutes. (Note that award years refer to when the candidate movies were released and not when the event occurred. For example, the 90th Academy Awards event was held in early 2018 for movies released in 2017, so in this piece we will refer to it as the 2017 awards.) But to what extent has the Oscars grown over the years, both in runtime and in length of speeches? Thanks to the Academy Awards Acceptance Speech Database (aaspeechesdb.oscars.org), and some information from Wikipedia, we can let the data speak for themselves.

Our data set spans 75 years, from 1942 to 2017, with some missing entries for 1943–1945. Oscars were awarded for those years, but the speech transcripts are not recorded in the database. Speeches are currently being transcribed in reverse chronological order by the Margaret Herrick Library, a Beverly Hills library dedicated to the history of the motion picture, which has complete transcripts from 1956 and piecemeal entries prior to that year. Ceremony runtimes were collected from Wikipedia (bit.ly/2Uz7c7W). For individual speeches, runtimes were not readily available, so word-counts were considered instead.

In this article, we will look at speeches in three of the most iconic categories: best actor in a leading role; best actress in a leading role; and best picture. The total word-count for each speech was tabulated only for the specific winners – or whoever was accepting on the winner's behalf – with extraneous banter removed. The most notable case of this was the 2016 best picture debacle, when *La La Land* was incorrectly called up prior to *Moonlight*, the true winner; our word-count begins when *Moonlight* producer Adele Romanski and director Barry Jenkins take the mic.

Besides scrubbing such banter, a few outliers were removed from the data before analysis: these are Fred Zinnemann's 1966 best picture and William Holden's 1953 leading actor speeches, each consisting of just four words, as well as Katharine Hepburn's 1981 leading actress Oscar, for which she was not present.

Times are changing

Beginning with ceremony runtime in minutes, a cursory plot of the data will reveal a noticeable changepoint at the 1972/73 boundary, which is during the brief period in the early 1970s when NBC took over the Oscars television broadcast from ABC. When considering such a data sequence, a changepoint is a point where some aspect of the data distribution suddenly changes, such as the mean value increasing drastically. For example, the mean value of runtimes for 1942–1972 is 133 minutes, whereas the mean value of the runtimes for 1973–2017 is 212 minutes.

To make this more precise, we use the `changepoint` library in R.[1] The `cpt.meanvar` routine searches for points in the data where the mean and variance change significantly with respect to some penalty. In our example, the mean increases quite abruptly at 1972/73 while the variance suddenly drops. In order to avoid overfitting a model with many such points, a penalty

term is applied. This can be thought of as adding a cost for each additional changepoint included in the model. Hence, only the most drastic changes are included in the final model to justify the cost. Multiple segmentation methods can be implemented with a variety of penalties. Applying either the pruned exact linear time (PELT) method or binary segmentation with the modified Bayesian information criterion (MBIC) penalty (see box, page 26) finds the single obvious changepoint between 1972 and 1973. Subsequently, modelling the runtimes as a linear function of the year with a single jump at 1972/73 gives an estimated jump of 55.8 minutes and a small but nonetheless statistically significant average increase of 38 seconds per year (Figure 1).
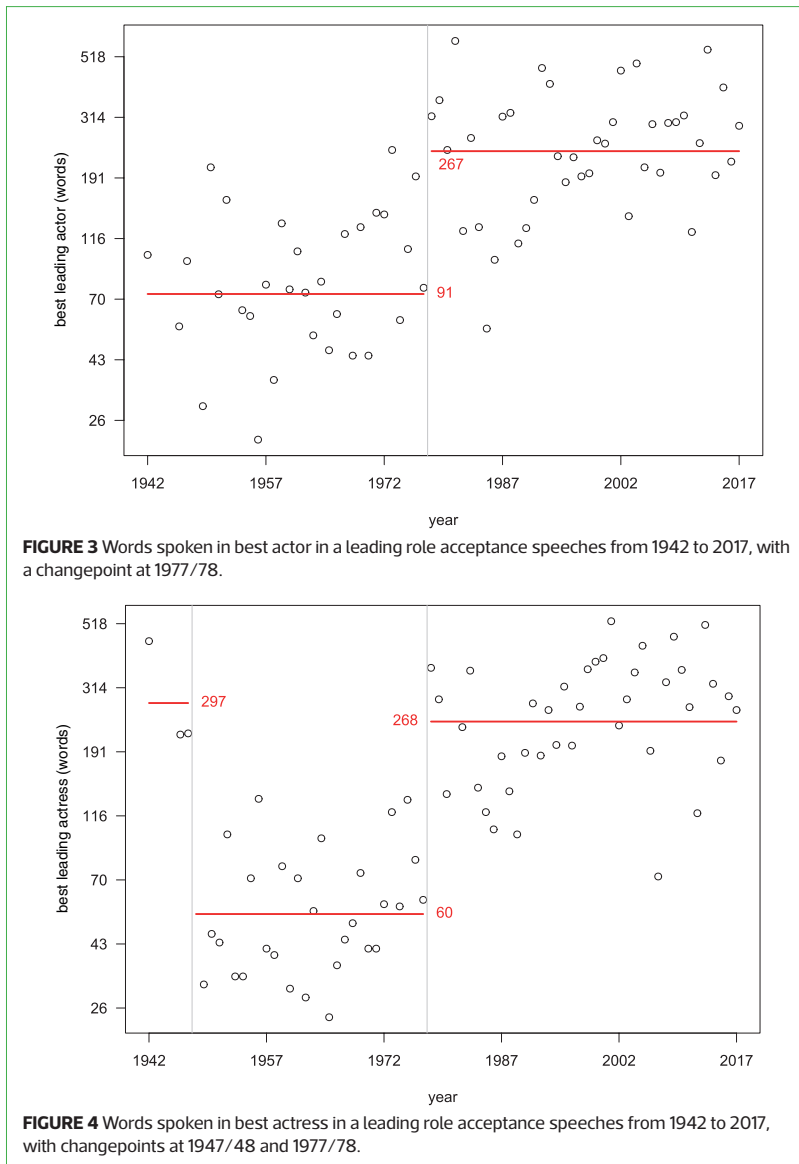
A single explanation for this increase in runtimes is unlikely. Discussing this result with a reference librarian at the Margaret ▶

FIGURE 1 Runtime in minutes of Academy Awards ceremonies from 1942 to 2017, with a changepoint at 1972/73 and an average annual increase of 38 seconds.

FIGURE 2 Words spoken in best picture acceptance speeches from 1942 to 2017, with a changepoint at 1978/79.

FIGURE 3 Words spoken in best actor in a leading role acceptance speeches from 1942 to 2017, with a changepoint at 1977/78.

FIGURE 4 Words spoken in best actress in a leading role acceptance speeches from 1942 to 2017, with changepoints at 1947/48 and 1977/78.

Herrick Library, I was told that the specific awards producer and host, the number of awards given, and any musical numbers or special anniversary events all have a strong influence on ceremony length. And there may be other factors to consider – not least the terms of the contract between the Academy and its chosen broadcaster, which govern such things as the timing of commercial breaks and conditions for whether or not a broadcast can run over time. (In 1950, for example, many award winners were not given the chance to speak due to severe time limits on the live radio broadcast.) Therefore, the switch from ABC to NBC and back again in the 1970s – presumably accompanied by new contracts – may have contributed to the jump in runtimes. Indeed, TV broadcasters may have favoured a longer ceremony to allow for more commercial breaks, as these can be quite lucrative.

Details on MBIC

When looking for changepoints in the runtime data (Figure 1), using the unmodified BIC identifies an additional changepoint at 1993/94. However, this is most likely an artefact of the extreme broadcasts for the 1998, 1999, and 2001 awards, which all ran over 4 hours. More extreme is the use of the Akaike information criterion (AIC), which liberally identifies 17 changepoints over the 75-point data set. The `changepoint` library in R defaults to the MBIC penalty with good reason: developed by Zhang and Siegmund, the MBIC penalty is a modification of the classic BIC specifically for changepoint detection.[2] While the classic BIC, given a sample size of n, imposes a penalty of $\log(n)$ for each changepoint included in the model, the MBIC increases this to $1.5 \log(n)$ and includes a second term summing the logarithms of the proportion of datapoints lying in each interval. This second term is maximised when the changepoints are equally spaced, and though dominated by the first term, still acts to counterbalance it. In the Oscars ceremony runtime data set, a single changepoint would result in a penalty of $1.5 \log(75) = 6.5$ from the first term. If the changepoint were at the 37th datapoint, then the second term will reduce the total penalty by 0.7. If this changepoint were at the third datapoint, then the reduction would be 1.6. Thus, equally spaced changepoints are more harshly penalised (i.e. cost more to include) than those that are clumped.

According to trade magazine *Variety*, ABC wanted between $2m and $3m for each 30-second ad slot during this year's Oscars broadcast (bit.ly/2PK8GZm). However, with the show's TV audience reportedly shrinking in recent years, it may be that viewers are not so keen on longer runtimes, and there are reports that the format of the show will be changed this year to try to keep the ceremony to 3 hours (bit.ly/2PPvyXu).

I next applied the same methods of analysis to the log of the total word-count of acceptance speeches over the same years for the categories of best picture, actor in a leading role, and actress in a leading role. Changepoints were discovered at 1978/79 for best picture (Figure 2, page 25) and at 1977/78 for the actor and actress awards (Figures 3 and 4). There were no statistically significant rates of increase over the years, except for leading actress speeches, whose total word-count is increasing by about 1% per year. Furthermore, a changepoint at 1947/48 was detected for the leading actress speeches. While the actress data set only contained three speeches prior to 1948, one of those was the 1942 award given to Greer Garson, who spoke for nearly 4 minutes. Her speech was not fully recorded, so its transcribed length of 452 words is only a lower bound on the total. The apocryphal story is that this lengthy speech is the reason time limits were imposed.

Counting thanks

No matter how many words are spoken, the common theme in acceptance speeches is "thank you". Thank you to co-stars and family members, producers and agents, industry bigwigs and the "little people".

Across all the years of our data set and the three types of speeches – actor, actress, picture – the number of "thanks" and "thank yous" was recorded. To model this count data, a Poisson regression was used with the log-link function. The Poisson distribution is used to model counts of rare events, like lightning strikes or radioactive particles hitting a detector; the Poisson regression assumes a parametric model where the number of "thanks" follows a Poisson distribution with rate parameter equal to the log of a linear combination of the inputs or predictors. Three predictors for the number of "thanks" were considered – the year, the total word-count, and the type of speech – giving the following equation:

$$\log(E(thanks)) = \beta_0 + \beta_1(year) + \beta_2(words) + \beta_3(typeActor) + \beta_4(typeActress)$$

where $typeActor$ and $typeActress$ are dummy variables, making β_3 the difference between best picture and best actor "thanks", and β_4 the difference between best picture and best actress "thanks". The best picture category is treated as a baseline and does not explicitly appear in the above equation.

Based on this model, there was no significant difference in the number of "thanks" across the three types of speeches considered. However, for every successive year, there was an increase in "thanks" of about 1.4%, with 95% confidence interval [1.0%, 1.9%]. There was also a statistically significant increase in "thanks" of 0.23%, with 95% confidence interval [0.17%, 0.29%], for every additional word spoken. This provides evidence that longer speeches are in fact longer because there are ever more people who require thanks, and each year brings with it more people to thank.

The Poisson regression identifies a few anomalous speeches in the data set. At the high end, there is Halle Berry's 2001 leading actress Oscar speech for *Monster's Ball*, in which 32 of her 528 words were "thanks", and Rod Steiger's 1967 leading actor Oscar speech for *In the Heat of the Night*, where about one in every 11 words was "thanks". At the other end of the spectrum, Tom Hanks's 1993 leading actor Oscar speech for *Philadelphia* ran for 414 words without a single "thanks", though it is worth noting that the speech still conveyed his gratefulness to his co-stars and friends. Hence, while the word "thanks" is perhaps the most straightforward way of showing gratitude, a true measure of a speech's thankfulness goes beyond merely counting "thanks".

Jimmy Kimmel and the jet ski

Finally, let us return to the story of Jimmy Kimmel and the jet ski. Comedian and talk show host Kimmel hosted the Oscars for the years 2016 and 2017. After nearly 4 hours of awards for 2016, it seemed that the usual tactics of mic cuts and music cues were failing to temper the ever-lengthening speeches. So, Kimmel devised a plan for the 2017 awards to offer a jet ski to the person with the shortest acceptance speech (timewise).

A quick look at the data shows, disappointingly, that the total runtime for 2017 actually increased by 4 minutes over 2016. But was there a significant change in speech word-counts when compared pairwise over all the transcribed 2016 and 2017 speeches in the database? It appears not. A one-sided pairwise *t*-test comparing 24 speeches shows no significant change (*p*-value 0.26). On average, the speeches decreased by only 7.6 words year-on-year.

While Kimmel's plan failed to curb speech lengths, it does motivate an intriguing but expensive follow-up study: how many jet skis are required to cause a statistically significant drop in speech lengths? ∎

References
1. Killick, R. and Eckley, I. (2014) changepoint: An R package for changepoint analysis. *Journal of Statistical Software*, **58**(3), 1–19.
2. Zhang, N. R. and Siegmund, D. O. (2007) A modified Bayes information criterion with applications to the analysis of comparative genomic hybridization data. *Biometrics*, **63**(1), 22–32.