80 Years of Popular Music

A quantitative analysis of pop music trends

Sepehr Sobhani
University of Mary Washington
Supervised by: Dr.Mark Snyder , Dr.Theresa Steward

May 3, 2016

Abstract

The music industry nets 15 billion dollars a year, and unlike many other industries, is largely immune to market fluctuations. If we are to believe that people vote with their money, the average American must value music enough to sacrifice other aspects of his/her life to avoid a decrease in spending on music. Music is also an expense that spans individuals across the economic spectrum. This makes research into pop music relevant to a majority of people. This paper will use recent advances in the availability of data on music to draw conclusions on how musical tastes change in the face of changing economic, political and social times.

1 Introduction

1.1 Background

There is wealth of data on the internet specifically about music. This data ranges from subjective critics to objective metrics on songs. This research will focus on the latter. By synthesizing the top songs in the U.S by year since the 1940s, and combining the data with meta data on individual songs, a conclusion for how the industry has changed over the years, and how it reacts to change can be discerned. Research into pop music is not a new phenomenon. Because of how lucrative a song has the potential to be, research is often done on what makes a pop song successful, or what is fashionable for a given month or week in pop music. The significance of the research being proposed, and what sets this research apart from previous projects is this study will cross correlate aspects of music with other national metrics and political events. Ideally, conclusions drawn from this study will allow music to be used as an economic and political indicator to aide in public policy decisions. This study will also be holistic and look to draw conclusions on trends spanning 80 years, a time frame that is larger than most other studies. The research will demonstrate that music is more than simple source of entertainment and can work as a unifying indicator of the national psyche.

1.2 Weighted vs. Unweighted Results

During data capture, for every week of any given year, a range between 10 and 150 songs were gathered for analysis. This means that a song could and often does appear week after week for a certain period of time, so there are two ways to analyze the data. Either a song can be given extra weight for the number of times it appears throughout the year, more accurately representing its importance in the public psyche, or the song can be counted only once, allowing for more consistent trends. Most charts in the paper will include both a weighted and unweighted analysis of the data.

1.3 History of the top charts

On January 4th, 1936, Billboard magazine published its first top music charts. At first, the Billboard only aggregated the top 10 songs in the country by surveying music merchants across the country. In mid-1958, Billboard consolidated data and started publishing the top 100 songs in the country using data about sales, airplay and at the time, jukebox activity. Currently Billboard uses sales, radio airplay, digital downloads, and streaming activity (including data from YouTube and other video sites) to get an accurate representation for how popular a song currently is. It is important to keep in mind how much more democratic the system of the top charts has become. In 1950 the average record cost $38.85 adjusted for inflation. This is on top of the already expensive record player required to play one. This meant that most people would not be able to "vote" by buying music and affecting the charts. With the advent of technology, music has become more ubiquitous, allowing people to "vote" for their songs by simply clicking on a link. The Top Charts, now more than ever,

are an accurate representation of peoples musical tastes spanning across the socioeconomic strata.

2 Data Collection

A custom web crawler was written using C# and crawled "Ultimate Music Database" (http://www.umdmusic.com/) using a web scraping technique. The data from UMD contains where a given song was in the charts for a given week, the songs name, and the songs artist. All the data is stored in a normalized Microsoft SQL Database hosted locally. The database schema is shown in figure 1. The songs were then enriched by making API requests to the Echo Nest, Musixmatch and Microsoft cognitive services. Echo Nest Meta Data parameters include:

1. Acousticness
    - Represents the likelihood a recording was created by solely acoustic means such as voice and acoustic instruments as opposed to electronically such as with synthesized, amplified, or effected instruments. Tracks with low acousticness include electric guitars, distortion, synthesizers, auto-tuned vocals, and drum machines, whereas songs with orchestral instruments, acoustic guitars, unaltered voice, and natural drum kits will have acousticness values closer to 1.0.
2. Danceability
    - Describes how suitable a track is for dancing using a number of musical elements (the more suitable for dancing, the closer to 1.0 the value). The combination of musical elements that best characterize danceability include tempo, rhythm stability, beat strength, and overall regularity.
3. Duration
    - Length of the song in seconds
4. Energy
    - Represents a perceptual measure of intensity and powerful activity released throughout the track. Typical energetic tracks feel fast, loud, and noisy. For example, death metal has high energy, while a Bach prelude scores low on the scale. Perceptual features contributing to this attribute include dynamic range, perceived loudness, timbre, onset rate, and general entropy.
5. Key Signature
    - The key of the song
6. Liveness

- Detects the presence of an audience in the recording. The more confident that the track is live, the closer to 1.0 the attribute value. Due to the relatively small population of live tracks in the overall domain, the threshold for detecting liveness is higher than for speechiness. A value above 0.8 provides strong likelihood that the track is live. Values between 0.6 and 0.8 describe tracks that may or may not be live or contain simulated audience sounds at the beginning or end. Values below 0.6 most likely represent studio recordings.

7. Loudness
    - Average volume of the song in Db.
8. Mode
    - Shows if the song is major or minor
9. Tempo
    - Tempo of the song in BPM
10. Time Signature
    - Time signature of the piece
11. Valence
    - Describes the musical positiveness conveyed by a track. Tracks with high valence sound more positive (e.g., happy, cheerful, euphoric), while tracks with low valence sound more negative (e.g. sad, depressed, angry). This attribute in combination with energy is a strong indicator of acoustic mood, the general emotional qualities that may characterize the track's acoustics. Note that in the case of vocal music, lyrics may differ semantically from the perceived acoustic mood.

The EchoNest API is a well-regarded music meta data database that was recently acquired by Spotify Inc. In regards to how acoustic attributes are calculated, the Echo Nest states the following:

"An acoustic attribute is an estimated subjective quality of a track. It is modeled through learning and is given as a single floating point number ranging from 0.0 to 1.0."

After all the data has been archived, it is fed into Tableau, a business intelligence tool, for visualization synthesis and analysis.

3 Analysis

3.1 Insights into Tempo

Figure 1 shows the median tempo from 1940-2016. It can be observed that tempo has taken a parabolic trend over the past 80 years, with the peak occurring around 1981. The trend in tempo also tends to be smooth showing that changes in tempo preferences are gradual

and not sudden. The upward trend in tempo starts around 1955, coinciding with the rise in popularity of rock music as shown in figure 11. Decrease in tempo coincides with the decline in popularity of rock music, and the rise in popularity of R&B as shown in figure 11. This leads to the conclusion that tempo fluctuations in pop music in the 20th century can be largely attributed to rock music.

3.2 Insights into Volume

Figure 13 shows average decibels per year from 1940-1980. A bread overview shows that there has been a positive trend in the volume of music over the past 80 years. Upon closer inspection of the data, we can see a period of stagnation in the volume of music from 1960 to 1990. Before and after this period there were and continue to be dramatic increases in the volume of music. Using evidence from figure 19 we can determine that this trend is not a result of the volume of genres. Using rock music because of its longevity, we can observe that rock music has been increasing in volume. This same trend holds true for all genres. A more likely explanation is "The Loudness Wars" and figure 13 provide strong evidence to support the existence of this phenomenon.

The Loudness War is the name given to the increase gradual increase in loudness levels. The public colloquially uses the term "Loudness War" to refer to a recent (past 20 years) phenomenon of increasing loudness. Figure 13 show that the most recent increase in volume starts around 1991. This is the same year Compact Discs(CD) reached mainstream popularity. The use of CDs signified a departure from the analog medium in which music was being stored in, to the digital format we know today. Digital music format allows for an increase in dynamic range that analog formats do not allow for, and this can be seen in figure 13. More interestingly, an increase in decibels can be seen from 1940 to 1960 showing that the "Loudness Wars" are not a recent phenomenon. This 20-year span shows engineers quickly reaching the limit of 7" Vinyl records, a limit that would not be broken for the next 20 years. It is also important to note that the "Loudness Wars" of the past 20 years have not plateaued and music will continue to increase in average loudness for the near future.

3.3 Insights into Key

Figure 6 shows key share for any given year. Each zone represents a key, and is split into two sections by a line. The area above the line is minor for that key, and the area below is major for that key. The time series reveals numerous trends about the composition choices of musicians. The first conclusion drawn, is that key popularity after the death of jazz has become stable. Changes in key popularity happen gradually, mirroring the gradual changes in genre popularity, show in in figure 11. The starkest observation is the decrease in the popularity of the key of Bb major. Bb major is the best key for brass instruments. The decrease in popularity of Bb shows the abandonment of brass instruments in the first half on the 20th century. Figure 7 shows total counts for major keys for all samples. The popularity of the keys used follows a distinct pattern: C(no accidentals) , G(one sharp), D(two sharps), A(3 sharps), F(one flat) before the pattern is broken by the key of C#. This is because guitars have been a more popular instrument for composition than pianos. The popularity of C# can also be explained by the use of a capo, making the guitar play higher for a singer.

Figure 8, 9, and 10 all display minor modes, filtered with different timeframes. The data from these figures is far less conclusive then figure 7. One would expect A minor it be the most popular mode, being the relative Minor of C Major. This is not the case. Overall, A minor the 3rd most popular minor key, which is unexpected. To determine whether this is a result of the popularity of guitars in composition, data was split into two groups: 1988-2016(figure 9) and 1940-1970(figure 10). Figure 9 shows A minor drop to 7th in popularity, while figure 10 shows A minor rise to second, within a margin of error of F minor. These results are consistent with the decrease in popularity of rock music the first half of the 19th century. Overall, we can conclude that composers do gravitate to certain keys when composing, and this key is largely dependent on the instruments used in music being played.

3.4 Insights into Lyrical Sentiment/Valence

Sentiment analysis is a method of discerning weather a textual input is positive or negative in sentiment. Although the uses for sentiment analysis are usually limited to the business world (eg. Determining whether a product review is positive or negative) this project used Microsoft's recently released sentiment analyzer on music lyrics from 1980 to the present. Sentiment goes hand in hand with valence; a measure of how musically positive a song is, not including lyrics.

Figure 15 shows average sentiment for a given year in a time series. Attempts to use lyrical sentiment as an indicator for economic trends, or consumer confidence proved inconclusive. Lyrical sentiment does not correlate with the University of Michigan's consumer confidence index, oil prices, or unemployment. There are a number of explanations for this. One is that the sentiment analyzer was inaccurate because its main purpose was not for lyrics. Another is that the sample size was too small. Another is that a person's mood does not correlate with the kind of music they listen to. There is also the issue that the Billboard charts do not react quickly and songs slowly rise and fall. This means that the Billboard charts will not be responsive to emotional shocks. Figure 20 shows sentiment for every given month. After the September 11th terrorist attacks, lyrical sentiment falls for 5 consecutive months before reaching the second lowest sentiment for any given month. However, this could be attributed to random chance. John F. Kennedy's assassination in 1963 shows no large impact on lyrical sentiment. This could possibly be attributed to different methods in collecting Billboard 100 data, but is more likely random variation. More research must be conducted to draw a more divisive conclusion.

Figure 16 gives insight into weather more positive sounding music has more positive lyrics. Simply looking at the scatter plot allows one to conclude that there is no correlation between how positive a song sounds, and how positive the lyrics are.

Figure 21 shows valence by month. It can be observed that valence slowly rises, peaking in the summer, then declines in the fall and winter. This follows the idea that summer jams tend to be more positive. The same trend does not hold true for lyrical sentiment, further supporting the theory that lyrical sentiment and song valence are not correlated.

Finally, Figure 18 displays average lyrical sentiment by genre. Country music has a reputation for being largely and exclusively about heartbreak, and is often nostalgic, and graph displays country music as the most negative genre by a large margin.

3.5 Insights into Genre

Figure 11 shows the percentage share of a given genre for each year. This is an effective way of visualizing genre popularity over time. The chart allows pop music over the last 80 years to be effectively divided into 3 sections: 1940-1956(The death of jazz), 1956-1990(The reign of rock) and 1990-present (genre diversity).

The 1940-1956 timeframe holds the most interesting trends. First is the striking rise, and immediate death in the popularity disco music (1975-1979). Over the course of 3 years, disco rose from obscurity, to almost 20% genre share, then immediately going down to 0% share in one year. This is never seen again and can be attributed to the strong anti-disco movement, and can even be traced to "Disco Demolition Night" in 1979. The event is commonly attributed to the death of disco, and figure 11 supports this view. Secondly, figure 11 shows how popular rock music has been every since the advent of the Billboard 100, peaking at 55% percent of charts being rock in 1981. No other genre has received this level of popularity for such a long period of time. We can also see when rock music started to decline, in 1981. This is most likely attributed to the birth of MTV that same year. MTV became a taste maker after 1981, dictating the direction of music around the mid to late 90's. The rise of r&b also was a factor in the decline of rock music's popularity.

In the 24 years we have entered an age with more musical diversity than any time in the modern era. There is currently no dominant genre like there was in the 1980's, with all genres sharing roughly 20% share of the top charts, with nothing indicating a stop to this trend.

3.6 Insights into Genre Churn(Experimental)

In order to analyze how volatile the life and death of genres are, a new metric was developed in order to measure the volatility of the Billboard 100 charts in terms of genre. The following equation was used to measure genre volatility:

$$\sum |\Delta\theta| \tag{1}$$

Where θ is the percentage share of a genre, and summation is calculating the total across all genres. Δ uses the change is a genres share for a given year from the previous year. Figure 17 shows genre churn over time with larger numbers meaning larger changes from the previous period. The graph shows genre churn sharply declining after the death of jazz, and slowly declining even further with the rise of rock music(1960-1981). The decline in the popularity of rock music shows genre volatility rise until stability is reached after 2000. The chart concludes that musical genres come and go consistently and that even a period of high volatility is not far from average volatility.

4 Figures

Go to `http://www.ssobhani.com/wp-content/uploads/2016/05/80-years-popular-1.pdf`