Movements and Measures: Developments in Music and Musical Influence

2125654

Contents

1 Introduction

Objective notions regarding music tend to be evasive. The language that non-experts use to describe it tends to be broad and highly variable, and the interpretation of such language is shaped by personal perspective. Describing the “feel” of a song can be reminiscent of describing the taste of a wine; a conversation partner may have a general idea of what “smoky” means, but they are likely to smile and nod rather than pursue a more specific understanding of the term. When tracking changes in artists, genres, and the state of music as a whole, more precise metrics provide a way forward.

Every musical artist has influences, and every musical artist has an impact on the music world, however small. Using the data sets provided, collected via Spotify’s API, we examine this network of influence, using metrics such as energy, danceability, and acousticness (among others) to track chains of influence, as well as trends in genre substance and popularity over time. We use this network of artists, its subnetworks, and various analyses based on given quantitative metrics in our investigation.

1.1 Metrics Present

Some metrics present in the song and artist data are more musically precise than others. Tempo, mode and key are common musical metrics, present in introductory musical vocabulary [1]. Loudness, using decibels, is straightforward. However, we also consider some less standard metrics that are present in the data set. We use danceability, a measure of how suited a song is for dancing; energy, which refers to the intensity or activity present in the song; and valence, which refers to the overall positive or negative feel of the song. Acousticness refers to the presence or lack of technological enhancement. Instrumentalness refers to the presence or lack of vocals, while speechiness refers to that of spoken words. Liveness is a measure of whether or not an audience is detected. In broader analyses, we also incorporate songs’ popularity. We analyze these various metrics provided by the data in order to better understand how music changes across time, genre, and through influence.

2 Assumptions

While our data set is not comprehensive, we assume in our analysis that it is a mostly comprehensive view of the music world. We note that this assumption suffers for the earliest and most recent years (such as the 2010s and the 1920s) due to less data present for these years. For this reason, there exist cases where data from these years may be somewhat misleading. We attempt to adjust our measures to account for this, whether that be via omission of data from these years, standardization in some form, or simply exercising caution when accepting conclusions based on these data. When any of these actions are taken, we describe them.

We also assume that the number of followers an artist has is an indication of how influential that artist was. This assumption may be problematic when analyzing the influence of more recent artists. Moreover, this assumption assumes that influencing one artist counts the same amount as influencing another artist, although intuitively we know that artists can have a different amount of influence on one artist than another. It also fails to account for chains of influence; however, we consider this factor at several points in our analysis.

Additionally, the data set contains some artists of unknown genre. We include these data points in our procedures, but note that they exist in a very small number, and therefore do not give any significant weight to properties and trends they exhibit.

3 Methods

3.1 Subnetworks

In order to observe the influence of the influencers on followers, we require a method of easily collecting the followers of a given influencer. Hence, using the artist influence data, we generate a directed graph in R using an adjacency matrix. With this adjacency matrix, we can perform operations such as applying similarity comparisons between influencers and followers, and generating subnetworks.

In order to generate relevant subnetworks of the given data, we require an algorithm to generate all neighbors (or all subsequent "followers") of a given node ("influencer") and vice versa. For this, we implemented a breadth-first search traversal algorithm in R.

3.2 Metrics of Similarity

We primarily analyze similarity based on seven metrics, all of which are provided on scales from 0 to 1. These metrics are danceability, energy, valence, acousticness, instrumentalness, liveness, and speechiness. Throughout this paper, we refer to our seven chosen metrics collectively as "similarity metrics."

We choose to eliminate some metrics from our similarity analyses because they are categorical, and others because of scaling. Because of this, our similarity analyses may fail to recognize the influence of the other metrics recorded in the data set on similarity between artists.

Once deciding to consider seven metrics ranging from 0 to 1, we considered the intuition behind genre and song similarity. Consider genres as regions in space, with songs of a given genre falling in its region. Each of our seven metrics represents a dimension, with some genres or songs being close with regard to some metrics and far with regard to others. Thus, the vector of the means of a genre's metrics would define a single point in space, as would the vector of a song's metrics. From this, we define our measure of similarity between genres, and by extension, songs, as the Euclidean distance between points in this space.

Further analysis could reveal that some metrics should be weighted more than others when determining similarity; a sampling of either public or expert opinion could illustrate that a subset of the metrics provided do more to influence a listener's subjective "picture" of a song, on average. In our analysis, we leave metrics unweighted, assuming equal influence on perception of similarity.

4 Analysis

4.1 By Year

4.1.1 General Trends

In an effort to observe multiple metrics at once, we plot them together and observe their trends over time. We do not include measures that are categorical (key, mode) and data which is difficult to observe on the same plot as other measures (duration, loudness, and tempo). Finally, we choose to prune out popularity because of a lack of interpretability. We are left with our seven similarity metrics, each on a continuous scale from 0 and 1. The line graph for these seven measures across time is shown in Figure 1. The graph that follows it, Figure 2, shows the same metrics' averages, calculated as averages over all artists rather than averages over all songs. We focus primarily on the graph derived from song data, Figure 1.

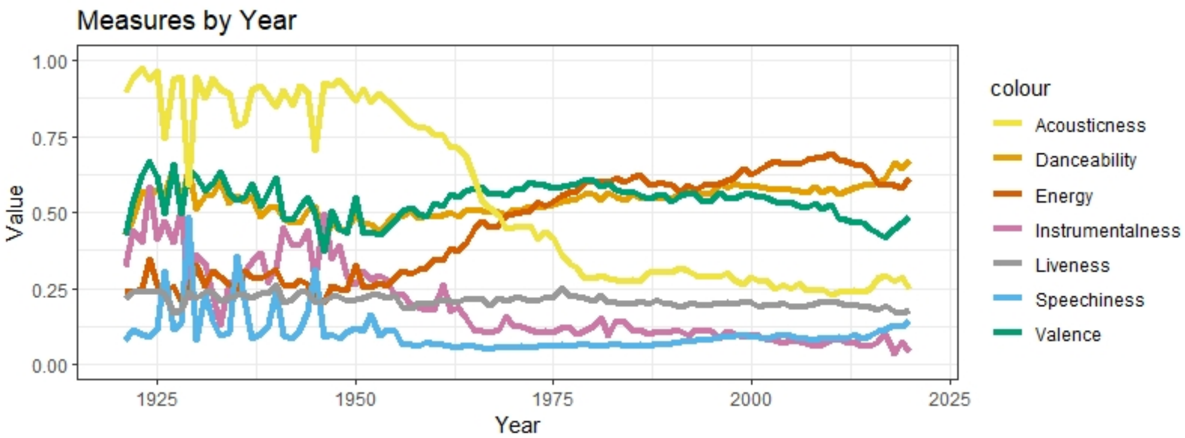

Figure 1: Average acousticness, danceability, energy, instrumentalness, liveness, speechiness, and valence of songs over time.

When examining the graph derived from song data, we see that there generally exists more variation from year to year for periods before 1950. After this period, there is a notably lesser presence of sharp peaks. We also observe some general trends. For the acousticness measure, we see that acousticness is generally high before 1950 at an average value of about 0.875, steadily decreasing from 1950 to 1975, and settling again at a value of around 0.33 after 1975. The downward trend continues, in line with continued technological developments. On the other hand, we observe that energy is generally low before 1950 at a value of 0.25, steadily rising until about 1985, after which point energy remains in a range of around 0.55 to 0.7, peaking at around 2010.

Moreover, we notice that instrumentalness follows an up-down-up-down trend, first peaking at around 1925, troughing at around 1935, peaking again at around 1947, and then steadily decreasing until around 1965, at which point the measure stabilizes. Overall, we observe that acousticness, energy, and instrumentalness see the most change, respectively, and that the other measures (danceability, liveness, speechiness, and valence) are relatively static.

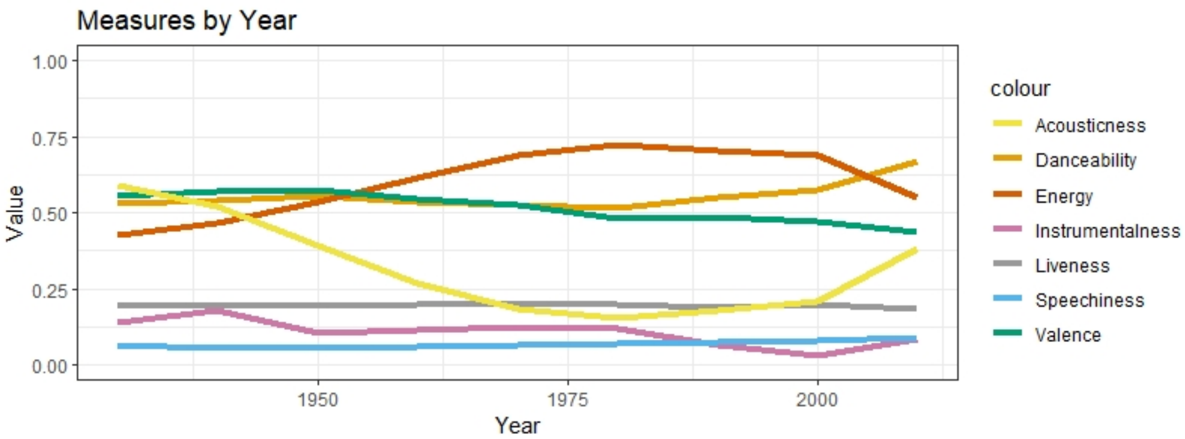

Figure 2: Average acousticness, danceability, energy, instrumentalness, liveness, speechiness, and valence of artists over time (active start date), weighted by influence.

Additionally, in Figure 2 we observe the same metrics as in Figure 1 over time, although instead of song data, we used average metrics of songs for each artists in a way that the contributions of each artist are weighted according to the number of other artists they influenced. We observe that the two graphs have noticeable similarities, suggesting that there could be a correlation between number of songs and influence of artist.

4.1.2 Metrics of Genres over Time

Figure 3 below illustrates to what degree a genre attracts new followers, and how that degree changes over time. We consider artists whose influencers are identified in the data set, determine their main genre, and plot the number of such artists who began releasing music in that genre by decade. Broadly, this paints a picture of which genres have been made popular by their respective influencers, and when these trends have occurred.[1]

[1]Since there exists significantly less data from years more recent than the 2000s, we omit data from these years in order to avoid suggesting that all genres are declining in popularity. As it stands, this illusion may still be created, albeit to a lesser extent, by data from the 2000s.

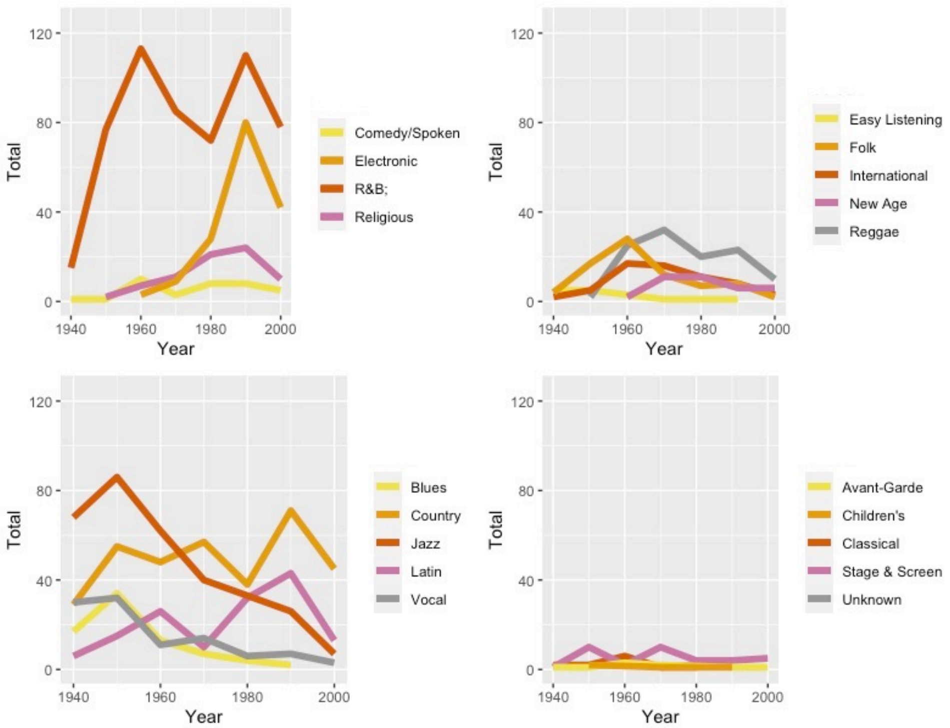

Figure 3: Number of followers given active start date, by genre. Data was divided into four plots for easier viewing. Plots exclude data for the pop/rock genre; scale would dwarf the other plots.

The graphs above in Figure 3 allow us to view simultaneous trends in genres from the 1940s to the 2000s. Some of the prominent trends appear to be the rise in R&B from the 1940s to the 1960s, the rise and fall of electronic, peaking in the 1990s, and the fall of jazz from the 1950s.

These graphs also serve to illustrate the high variability in popularity between genres overall. Even at their peak, many genres are far from nearing the popularity of others. This trend may be heightened by the data available in Spotify's API; it is possible that lesser popular genres have a lower proportion of songs present on Spotify, thereby compounding their seeming lack of popularity.

We also examine how genres have varied in content over time by averaging the measures of followers classified within that genre. Values over time of our seven similarity metrics are shown on the next page in Figure 4 for six genres. Pop/rock, R&B, electronic, and country are more popular genres with relatively high energy, while vocal and folk are less popular and less energetic on average.

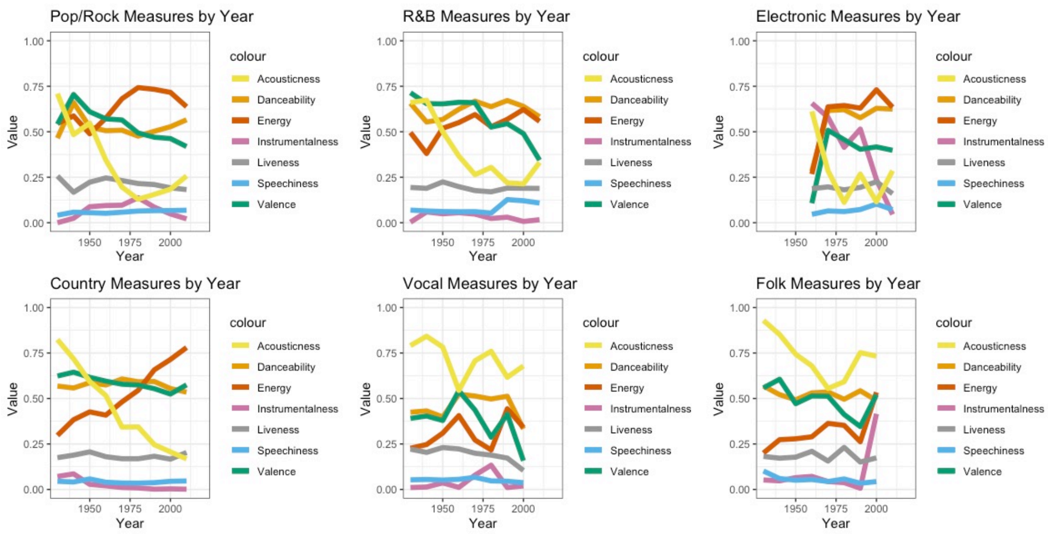

Figure 4: Average acousticness, danceability, energy, instrumentalness, liveness, speechiness, and valence of songs over time. Data from 6 genres ranging in popularity.

We see that, by and large, most metrics tend to remain close to constant over time, with some noteworthy exceptions. Acousticness trends starkly downward for pop/rock, R&B, and country, a trend that is present in some other genres not displayed in Figure 4. The same trend exists in the global song data, shown previously in Figure 1. It appears that for genres more musically tied to acousticness, such as vocal or folk, there is less of a downward trend. We also observe a strong increase in the energy of country music, which we suggest is influenced by a blending of pop/rock and country music in more recent times. We see a strong decline in the instrumentalness of electronic music, which may reflect a shift from a genre centered around new, electronic instrumentation to a genre of dance music with lyrics in the present.

We observe specifically that the energy of pop/rock, R&B, and electronic have only undergone minimal fluctuations between the 1960s and the 2010s, and that the same holds for folk and vocal music. This phenomenon is also present is most other genres, though there are exceptions like country music as mentioned previously. One interesting inference follows from this information. In Figures 1 and 2, we saw that energy across genres trended upwards between 1950 and 1985. Given the minimal increase in energy present in most genres (Figure 4) and the trends seen in genre popularity (Figure 3), we infer that the global upward trend in energy more reflects a rise in popularity of high-energy genres, rather than an increase in energy within genres.

4.2 By Genre

While analysis by year has a largely inter-genre focus, we also conducted intra-genre analyses. These include genre predictions, descriptive procedures regarding metrics within genres, and a specific dive into the pop/rock genre. We examine changes in pop/rock music's metrics over time, and choose a specific period in pop/rock to investigate how the influence of an artist correlates with their energy and when they began releasing music.

4.2.1 Classifying Songs by Genre Using Neural Networks

To examine what defines genres, we created a neural network that classifies whether or not the predominant genre of an artist who started in the 1990's was pop/rock using danceability, energy, valence, tempo, loudness, acousticness, instrumentalness, liveness, and speechiness. Through trial and error - in an effort to maximize predictive power - we created this neural network with a single hidden layer and 4 neurons in the hidden layer.

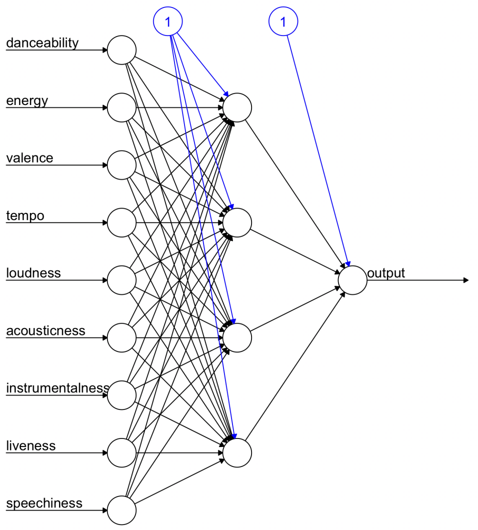

Figure 5: Representation of Neural Network

With this network, we were able to achieve an average of a 79.5% success rate over 20 networks with different test data in identifying whether or not an artist with a start date in the 1990's was a pop/rock artist or not. This suggests that there are fundamental differences between genres based on the quantifiable factors used in this neural net, especially when we control for a time period to ensure that characteristics of each genre are consistent.

However, while some genres seem to have defining characteristics, others appear very closely related. This can be seen by attempting to classify data from other genres using the same neural network. The genre of R&B, for example, can be separated fairly accurately from the pop/rock data. When all R&B artists from the 1990's are classified using the neural network, only 8.2% are mistakenly classified as pop/rock. However, country is not as easily distinguished. The neural network mistakenly classifies country artists from the 1990's 53.5% of the time as pop/rock artists. This may be because pop/rock and country in the 1990's share many characteristics due to a blending of the genres, as suggested earlier. Pop/rock and R&B, on the other hand, have fewer similarities, so they can be more easily sorted based on the metrics passed into this neural network.

4.2.2 Metrics of Songs by Genre

We analyze the distinguishing factors of genres, the similarities and influences between and within genres, and the relation between genres. We do this using the seven similarity metrics discussed previously, as well as tempo, loudness, and the proportions of modes of songs in a genre. Information regarding the genre of individual songs is not provided; however, we have data that includes influencers in music and their followers, including the predominant genre of each influencer and follower. If a genre is the predominant genre of an artist, then their music is predominantly of that genre. Therefore, we assume that all of their songs were of that genre in order to analyze broad trends and relationships in the characteristics of each genre.

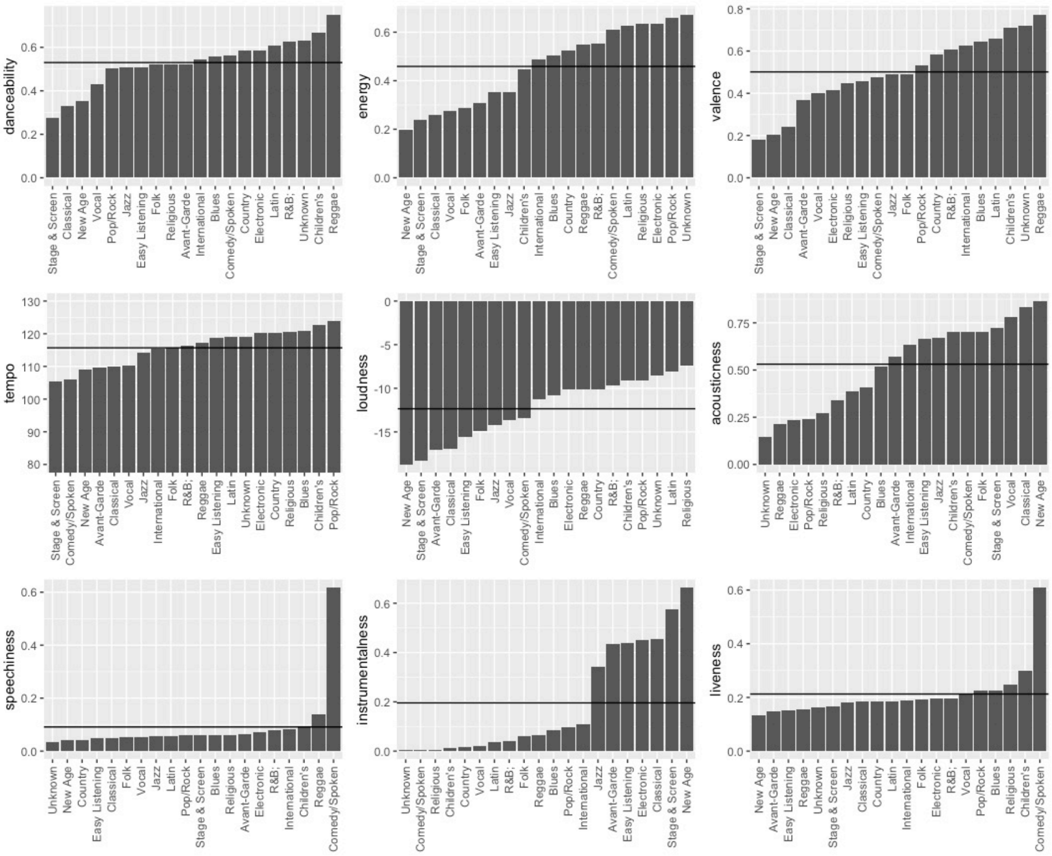

Figure 6: Average acousticness, danceability, energy, instrumentalness, liveness, loudness, speechiness, tempo, and valence of songs, ascending by genre.

Figure 6 above shows the genre mean of each of the following numeric metrics: danceability, energy, valence, tempo, loudness, acousticness, speechiness, instrumentalness, and liveness. Each horizontal line indicates the global mean of a metric.

It becomes apparent that many features are relatively consistent across genres, while others exhibit significant variability. Some metrics, such as liveness and speechiness, remain close to constant across most genres, with only select genres deviating. Others metrics exhibit more deviation between genres. Energy is perhaps the most significant example; we see a large difference between the energy of electronic and pop/rock music (hovering around 0.6) and genres such as vocal and folk (closer to 0.3). Differences like these, when combined with insights gained from studying changes in genre metrics and genre popularities over time, allow us to propose explanations for more global music trends in later analysis.

We also analyze one binary value across genres, illustrated below in Figure 7. The mode of a song can either be major, represented by a 1, or minor, represented by a 0. In Figure 7, we observe the proportion of all songs in a given genre that are in a major mode. The horizontal line indicates the global proportion across genres.

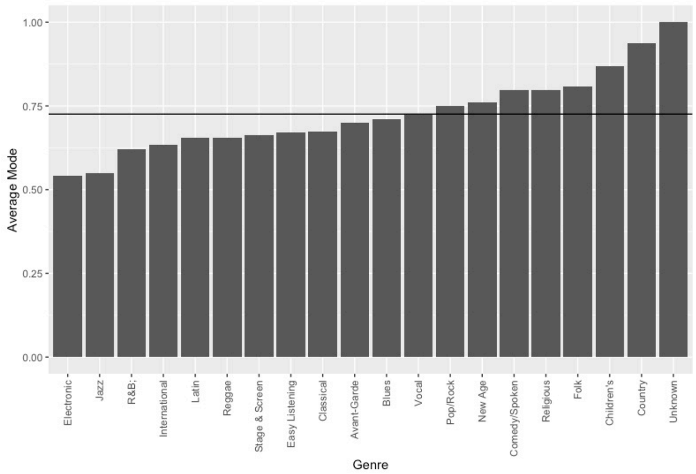

Figure 7: Proportion of songs in a major mode, ascending by genre.

The "average mode" of vocal music, which is 0.75, expresses that 0.75 of songs in the vocal genre have a mode of 1 and 0.25 of songs in the vocal genre have a mode of 0. While all genres have an average mode greater than 0.5, there is still variation between genres. Electronic and jazz both share the lowest mean modes with barely more than half of their songs having a major mode. On the other hand, country and children's have the two highest mean modes with the greatest share of their songs with a modality of major. Of course, we note that the mode of a comedy/spoken song has little to no interpretation.

4.2.3 Genre Analysis: Pop/Rock

Since we notice that the pop/rock genre has the most popularity of all genres, as observed in Section 4.1.2, we choose to study this genre in more detail. We investigate how influence data might paint the genre data in an alternative light.

From the graphs in Figure 3, we saw that over time, songs marked as pop/rock saw a downward trend in acousticness and valence, and an upward trend in energy over time. We might ask ourselves how this compares to the artist data (the averages of data over all of an artist's songs) and how this data differs when we simultaneously consider the influence of artists. In accordance with our earlier assumptions, the metric of influence considered is an artist's number of followers.[2] The unweighted and weighted graphs are displayed, respectively, in Figures 8 and 9 below.

[2]Because of we wish to consider influence, the graphs take data from the set of influencers, which is a large subset of all artists. It may also be worth noting that the "time" data refers to the decade of active start of the artist, and not the peak decade of popularity of the artist, which could potentially skew the results of the data.

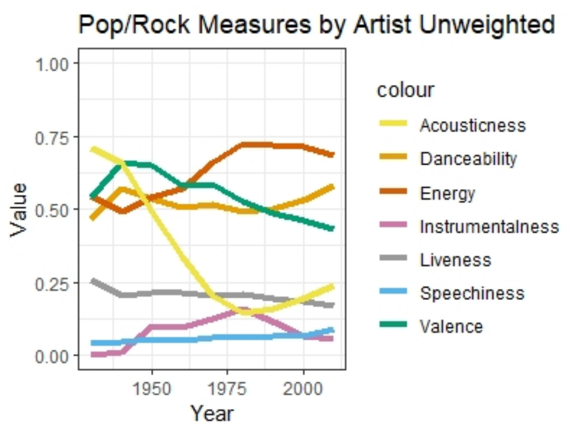

Figure 8: Mean over time of each metric for pop/rock influencers, unweighted.

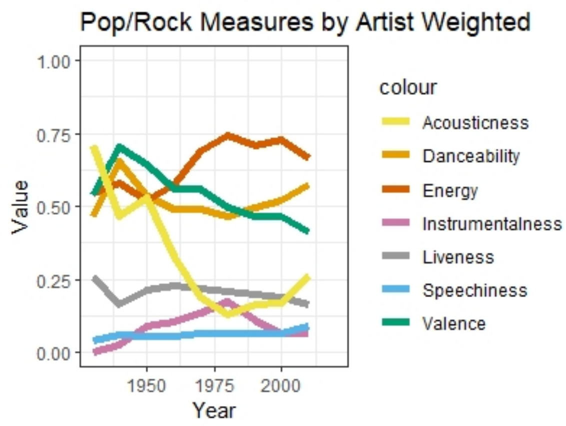

Figure 9: Mean over time of each metric among pop/rock influencers, weighted by artist's number of followers.

We may first choose to look at these graphs separately. As observed in Figures 4 and 6, we see that in both graphs, danceability, liveness, and speechiness are generally constant over time; meanwhile, we observe an upward trend in energy and a downward trend in acousticness.

In our comparison of the two graphs, we observe that they are rather similar, although we notice some distinctions; primarily, we observe that there are sharp peaks and troughs in the weighted data which are otherwise unseen in the unweighted data. This is reminiscent of the songs by year data in Figure 1, which show increased variance in data over time up until 1950. Some of these noticeable peaks and troughs include a drop in acousticness and liveness, and a peak in danceability, energy, and valence, in the year 1940. Also, after 1950, we notice a brief trough in energy in the year 1990.

4.2.3.1 Pop/Rock in the 1990s

Upon investigation of energy in the 1990s, we find that although the total average energy value lies at 7.27, as can be seen on a the graph when $x = 0$, the energy of the "top" influencers is generally lower, the lowest around 7.0. While this is only a difference of 0.3, we might look to see if this signifies any particular change.

While looking at influencer data, we choose to look at average follower active start date to convey information about what periods influencers had influence on, by minimum number of influencers, and observe whether this is consistent with energy. In Figure 11 we see that similar to the energy data, for those with influence on more followers in the data set, they tend to influence bands with an earlier start date. In particular, for influencers with a minimum number of 10 followers, we see that the average active start date is 1990, meaning that the majority of those influenced became active during the same decade.

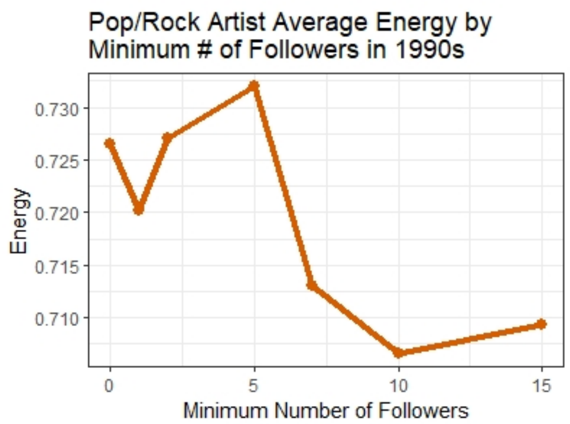

Figure 10: Average value of energy of the average song produced by artists, by minimum number of followers.

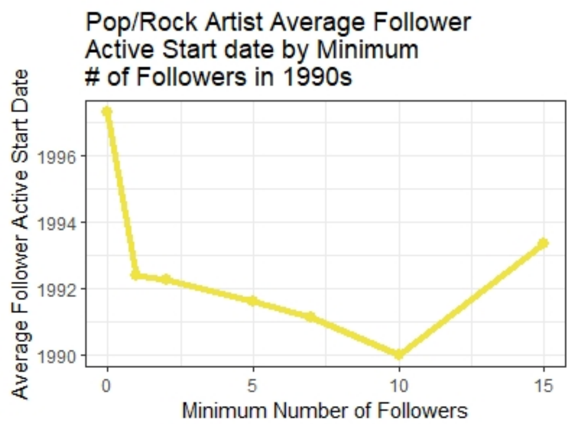

Figure 11: Average active start date of followers influenced by artists, by minimum number of followers.

This starkly contrasts the value considering the mean among all influencers, in which we observe a higher follower active start date. This could suggest that while more influential artists in the 1990s had lower energy songs on average, because their influence did not directly expand into the 2000s and 2010s, these artists took more influence from lesser known artists in the 1990s and/or from artists in their own times periods.

Additionally, with both of these plots, we notice that the data varies based on the number of minimum followers we consider. If we were to revisit Figures 8 and 9, Figures 10 and 11 would suggest that some of the data similarity between the plots may not be due to top influencers displaying the same characteristics as the rest, but rather that they may be similar due to other factors.

4.3 Influence

4.3.1 Influence by Year

We also examine how influence varies over time. Using the number of followers of an influencer as a metric for influence, we observe the number of influencers with 25-50 followers, 50-99 followers, and 99+ followers over time (according to active start date) and how these values compare to the total number of artists during that time frame.

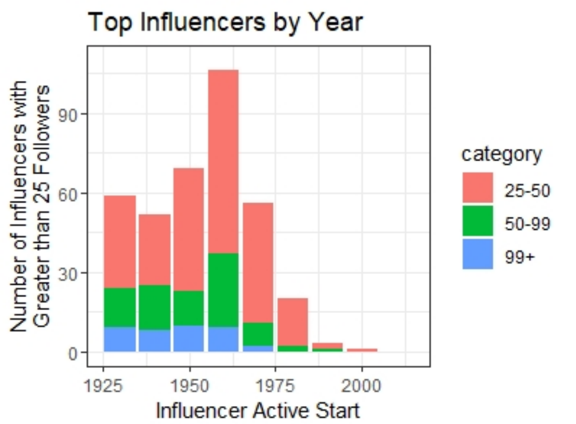

Figure 12: Followers of influencers with at least 25 followers, by active start date.

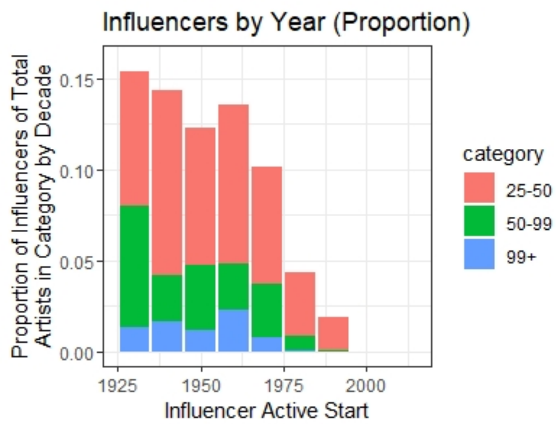

Figure 13: Proportion of influencers with at least 25 followers, by active start date.

As we see in Figure 12, the total number of influencers with 25+ followers is steady from the 1930s-1950s, peaks in the 1960s, then trends downwards. By the time we reach the 1980s, we're left with a few artists with 25-50 followers, and by the 1090s, few artists with 25-50 followers. This could suggest that over time, artists are taking influence from fewer sources or from a more dispersed set of sources. We also consider that artists from these more recent periods have yet to influence the artists of the future, so these totals are likely to increase with time.

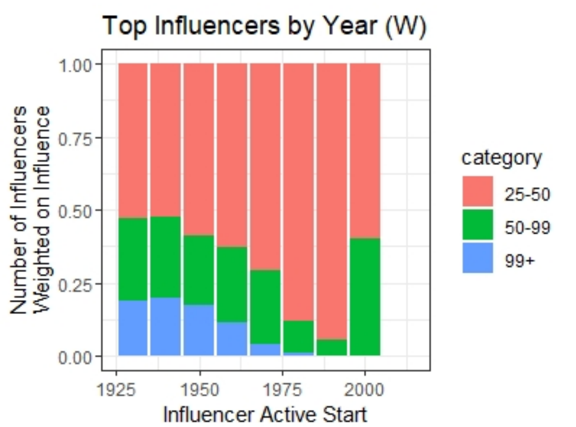

Figure 14: Proportion of influence provided by influencers with varying follower counts, by active start date.

Although the peak in the 1960s may appear to be outlier to such a trend, intuition suggests that the number of total artists grows as decade increases. This is illustrated by Figure 13, which shows the same data from Figure 12 but relative to the total artist count. Here, we see that the proportion of influencers with at least 25 followers generally remains constant until the 1960s, where it begins to trend down until the 2000s.

We notice several patterns when looking at the data conveyed in Figures 12 and 13. First, from the 1930s to the 1960s, the data suggest that top influencers are able to maintain the same relative influence over other artists, despite an increase in artists recorded in the database during this time. We also notice that from the 1960s on, influence trends down for all categories of influencers in both graphs, making the data difficult to interpret.

Finally, we turn to another graph to better interpret the data in Figure 13. Figure 14 shows the weighted influence of influencers with at least 25 followers according to decade. This means that the y-axis shows the proportion of influence (or proportion of number of followers) of influencers of the three classes. Here, we notice that until the 2000s, where a lack of data may impact results, the most influential group overall become the influencers with fewer followers (25-50) compared to the groups with more followers. Combined with the influencer count data in Figure 12, this again may lend itself to the hypothesis that over time, artists are taking influence from fewer sources or from a more dispersed set of sources.

4.3.2 Genres of Influencers versus Followers

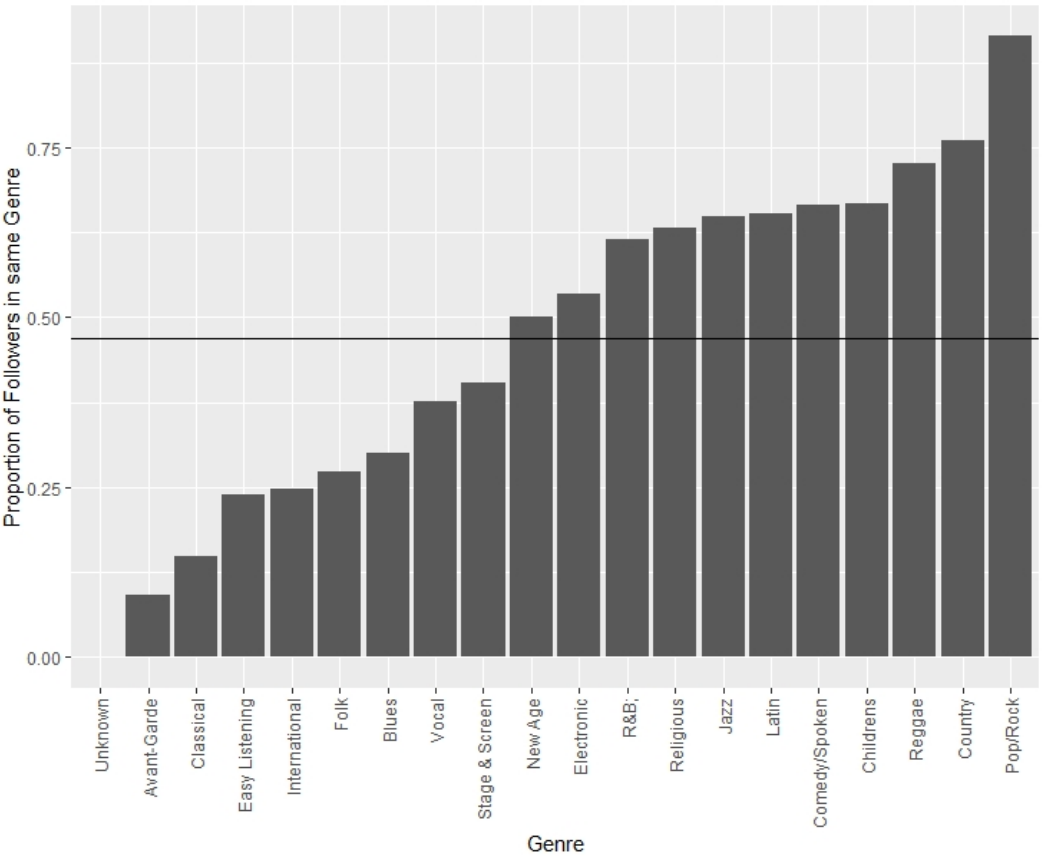

Figure 15: Proportion of followers with same genre as their influencers, by genre.

Figure 15 shows the proportion of influencers' followers in each genre whose music genre is predominantly the same as their influencers, and includes a line of the mean proportion across all genres. Out of $41,461$ influencer to follower relationships, $9,745$ of those relationships involve a follower with a different genre while $31,716$ involve followers with the same genre as their influencer. So approximately 76.5% of followers have the same genre as their influencer.

The amount of followers whose work is predominantly in the same genre as their influencers vary based on the genre of the influencers. Those influencers whose predominant genre is pop/rock mostly influence artists whose predominant genre is pop/rock. A similar, though weaker relation exists with country and reggae artists.

On the other hand, the data suggests that few artists whose predominant genre is avant-garde influence followers whose predominant genre is also avant-garde. We note that the same trend exists for classical and easy listening artists, which may seem counter-intuitive. It is possible that this trend exists simply due to a small sample size of avant-garde, classical, and easy listening artists, while many artists in more popular genres may still have taken influence from these artists.

4.3.2.1 Genres of Followers of Pop/Rock Influencers

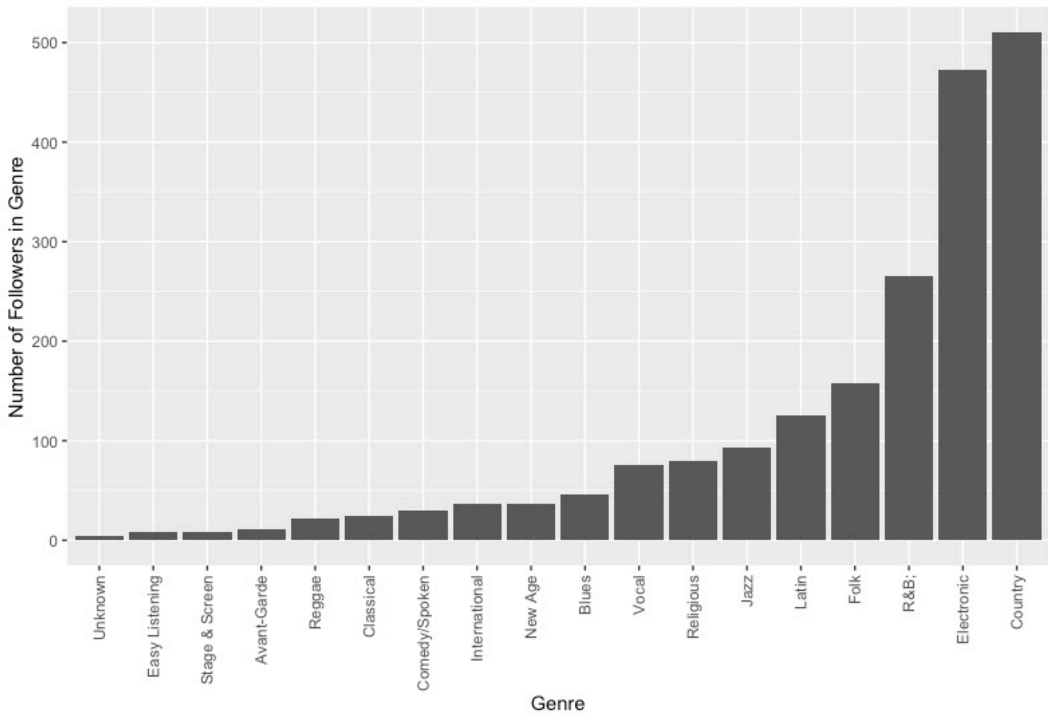

Figure 16: Followers of pop/rock influencers of a different genre, by genre.

Figure 16 displays the primary genres of followers of pop/rock influencers, excluding followers who are pop/rock artists themselves. We see that if a follower of a pop/rock artist does not make pop/music, they are very likely to make electronic or country music. Pop/rock influencers also influence a significant number of R&B, folk, and Latin artists. Pop/rock artists are least likely to influence artists whose predominant genres are easy listening, stage & screen, and avant-garde.

4.3.3 Subnetwork of Pop/Rock Artists in the 2000s-2010s

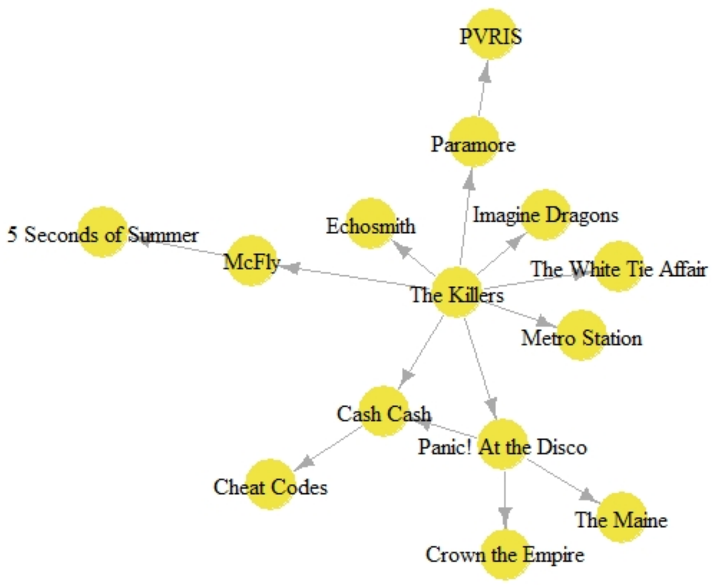

Figure 17: Subnetwork of followers generated from starting node "The Killers"

The above subnetwork in Figure 18 shows a chain of influence within the pop/rock genre. The 2000s band The Killers serves as the starting node, and the directional arrows illustrate which artists influenced which others. Using our similarity metric, we calculated the similarity between all pairs of artists in this subnetwork. We investigated whether pairs of artists directly connected in this subnetwork would exhibit higher similarity than pairs indirectly connected; we found a difference of virtually zero. We bear in mind that this may not be the case for other subnetworks based on other artists from different genres or time periods. It is possible that 2000s-2010s pop/rock is particularly homogeneous, or more specifically, the niche occupied by these fourteen artists is homogeneous. By sampling artists from different genres and time periods, then creating additional subnetworks, it would be possible to ascertain whether this phenomenon exists in other places.

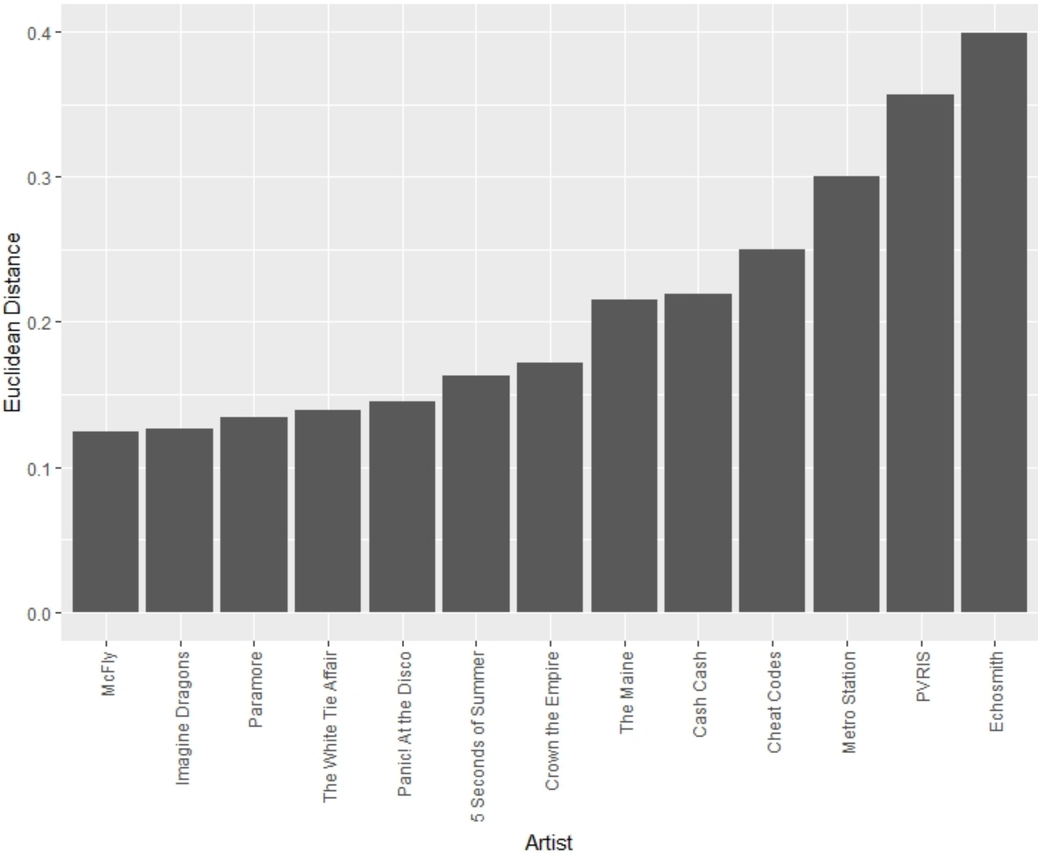

Figure 18: Euclidean distance from The Killers, ascending by artist

Figure 18 above illustrates the Euclidean distance between The Killers and each of their followers. We see that all followers are very similar to the Killers by our distance metric, so assessing relative rather than absolute difference is most appropriate. Even so, the maximum difference factor is less than 4x, reaffirming the homogeneity of the subnetwork.

4.3.4 Assessing Influence through Neural Networks

To examine the relationship between influencers and those who follow them, we created a neural network to predict whether or not a given artist is a follower of a given influencer based on the artists' average danceability, energy, valence, tempo, loudness, acousticness, instrumentalness, liveness, and speechiness. To improve our results, we restricted both the genre and time frame of our artists. In particular, we chose to examine whether or not we could predict if a given pop/rock band in the 1990's was a follower of the band Radiohead or not.

In our neural network, we used a single hidden layer with 3 nodes; trial and error revealed that this maximized predictive power. We fed in 90 teaching data points equally split between bands that were influenced by Radiohead and bands that were not. After creating 20 networks, we found our success rate to be 56.25% in predicting whether a given artist is influenced by Radiohead. The neural network's prediction accuracy was greater than that of random chance, but did not significantly improve predictions.

5 Discussion

5.1 Findings

We observed in Figure 1 that the values of the measurements for the average song change across the 1920s through the 2010s, particularly for acousticness and instrumentalness, which trend downward, and energy, which trends upward. Moreover, the graphs shown in Figure 3 allow us to observe how the number of followers of genres changed over time, in turn allowing us to understand shifts in the popularity of genres.

One benefit of studying these figures is assessment of the accuracy of cultural perception of trends. The downward trend in acousticness in Figure 1 appears to align with related technological developments; the advent of electronic music following World War II was marked by the manipulation of acoustic sounds using the new technology of magnetic audio tape [2]. We see many of these observed cultural trends in Figure 3 as well; for example, we see a rise in popularity of reggae during the 1960s and 1970s, in line with many people's idea of a "hippies with dreadlocks" era in American culture. We see that electronic music skyrocketed in popularity beginning in the 1980, which also tends to align with popular perception. However, some may be surprised to see that more new electronic artists emerged in the 1990s than the 1980s. Trends in individual genre metrics suggest ideas as well; the decline in instrumentalness of electronic music since its introduction suggests shifts in the defining nature of the genre, while the strong increase in the energy of country music would align with a blending between country and pop/rock.

Referencing our findings on global trends in metrics, intra-genre trends in metrics, and genre popularity over time, we were also able to provide an explanation for the general upward trend in energy across all music. By breaking up trends in popularity by genre, we find that genres that have become more popular over time are generally more high energy, and we rarely see significant change in genre energy over time. This suggests that the global rise in energy over time is due to increased prevalence and influence of already high-energy genres.

5.2 Limitations

Some limitations arise from the data present and from our model assumptions. Artist data given includes active start date, but not peak or active end date, meaning that some artists careers may be represented inaccurately. Additionally, the values of artist metrics present in the data (such as danceability, energy, etc.) are averages over all songs by the artists that are present in the data set; this overlooks the fact that certain songs by an artist may deviate significantly from their typical music, thereby pulling their averages away from a true representation of their work, and that different songs by an artist may have different levels of influence. We also note that since the genre of each song is not listed, only the primary genre of each artist, we adopted the assumption that all songs by a given artist are of that artist's main genre. This fails to account for the fact that many artists change genres or experiment with different genres throughout their careers.

Genre data for each song, or genre data for songs based on the genre of their album, would eliminate the need for this assumption.

We note that predictions based on our neural networks were better than chance, but had room to improve. We believe that with more comprehensive data, our prediction algorithms could be improved, thereby increasing their accuracy. This will be especially true for domains with less data, including the earliest and latest decades present, and underrepresented genres. As discussed briefly with its introduction, we acknowledge that our similarity metric is based on broad intuition, and could be made more appropriate through the implementation of data regarding public or expert opinion on what makes songs similar.

5.3 Future Directions

Improvement upon our calculation of similarity, from simple Euclidean distance to a more complex function, would likely yield more precise and nuanced results. An independent study aimed at gauging the importance of each metric in shaping individuals' perception of music would be useful in pursuing this goal. Data from such a study could reveal that certain metrics carry more weight when determining how similar people perceive songs to be, and our method of similarity calculation could be modified or replaced appropriately. This would require either a large-scale sample of the public, and subsequent analysis, or the input of professional musicians, music theory experts, and/or sonologists. We also believe that creating more standardized metrics for data such as tempo, loudness, and duration, and popularity would enrich similarity analysis.

Statistical assessment of the distributions of these measures is possible, and would be especially effective given a data set that is more comprehensive over all genres and time periods. This information would allow for appropriate standardization, allowing us to incorporate new metrics into our prediction algorithms. The additional dimensions would likely improve the accuracy and precision of our predictions, as well as our broader descriptive measures. We also propose that analyses of relationships between metrics may provide useful information as well. For example, strong relationships between certain metrics may suggest changes in what is included in similarity metrics, minimizing potential overfitting.

We also believe that the creation and implementation of additional subnetworks will give deeper insight into more specific intra-genre trends. For example, the homogeneity present in the subnetwork for The Killers may not be reflected in other subnetworks. Analyses such as these may reveal that certain genres are more variable, or that they tend to exhibit looser patterns of influence. These procedures could be conducted using the current data set with some additional time. We suggest drawing artists from a single genre, creating their subnetworks, and comparing their characteristics to our provided subnetworks and to each other.

References

[1] Western Michigan University. Glossary of musical terms. https://wmich.edu/mus-gened/mus150/Glossary.pdf.

[2] Stewart Wolpin. The race to video. *Invention Technology Magazine*, 10, 1994. https://web.archive.org/web/20110404045940/http://www.americanheritage.com/articles/magazine/it/1994/2/1994_2_52.shtml.