

# Genre and Temporal Dynamics in Spotify Popularity Prediction

Ryan Hausner

2026-03-31

The rise of music streaming platforms has created modern access to track-specific data, allowing for analysis of song popularity. While existing literature has explored audio features as predictors of popularity, less attention has been given to the combined role of genre and temporal dynamics — this paper addresses that gap. Using a public dataset of 114,000 tracks from 2000 to 2022, we apply a data science framework combining iterative OLS regression, interaction modeling, random forest, and rolling coefficient analysis to explore the predictive power of Spotify audio characteristics: loudness, danceability, energy, liveness, and valence, as well as genre and release year. Four iterative OLS regression models are developed using an 80/20 train/test split, showing that genre accounts for the largest gain in explained variation, increasing R-squared from 0.042 to 0.434. A genre-year interaction model further improves R-squared to 0.641, with interaction terms jointly significant confirmed by a partial F-test ( $F(80,318,243)=653.72, p<.001$ ), implying that the effect of genre on popularity varies across time — specifically that different genres rise and fall in prevalence at different periods. A random forest model confirms these findings, ranking genre and year significantly higher in feature importance based on impurity reduction. The most accurate model achieves RMSE=9.64 on a popularity scale of 0-100, with remaining variance likely attributable to unmeasured factors such as Spotify playlist algorithms and social media exposure. Rolling coefficient analysis further reveals the instability of audio features over time — energy’s contribution to popularity turned strongly negative post-2010, while danceability peaked around 2015-2016 — suggesting that the streaming era has fundamentally reshaped which acoustic properties drive popularity.

## Table of contents

**1 Introduction**

**2**

<b>2</b>	<b>Data Overview</b>	<b>4</b>
2.1	Summary of Dataset . . . . .	4
2.2	Popularity by Year . . . . .	4
2.3	Summary of Audio Features by Genre . . . . .	6
2.4	Popularity by Decade . . . . .	6
<b>3</b>	<b>Exploratory Data Analysis</b>	<b>7</b>
3.1	Popularity Distribution . . . . .	7
3.2	Popularity by Year . . . . .	8
3.3	Popularity by Genre . . . . .	8
<b>4</b>	<b>Model</b>	<b>9</b>
4.1	Baseline Model (Audio Features Only) . . . . .	9
4.2	Improved Model (Audio Features + Genre) . . . . .	10
4.3	Full Model (Audio Features + Genre + Year) . . . . .	11
4.4	Interaction Model (Genre x Year) . . . . .	12
4.5	Model Comparison . . . . .	12
4.6	Residual Analysis . . . . .	13
<b>5</b>	<b>Machine Learning Analysis</b>	<b>15</b>
5.1	Random Forest - audio features + year . . . . .	15
5.2	Random Forest Residual Analysis . . . . .	16
5.3	Model Comparison . . . . .	17
<b>6</b>	<b>Audio Feature Effects on Popularity Over Time</b>	<b>18</b>
<b>7</b>	<b>Results</b>	<b>20</b>
<b>8</b>	<b>Conclusion</b>	<b>21</b>
	<b>References</b>	<b>22</b>

## 1 Introduction

In the last two decades, the way music is consumed has shifted from physical media and radio to streaming platforms, fundamentally changing how listeners discover and engage with music. Platforms such as Spotify generate track-level data containing acoustic properties — measurable characteristics of a song’s sound such as danceability, energy, loudness, liveness, and valence — alongside a popularity index, a Spotify-assigned score from 0 to 100 reflecting recent stream counts and listener engagement. This creates an interesting statistical question:

how do acoustic features contribute to song popularity, and how does that relationship change across genre and time?

Popularity prediction provides insight for industry professionals and market researchers. Accurate popularity models could meaningfully inform production, marketing, and release decisions. However, previous models have struggled to explain a large portion of the variation in popularity, and critically, have not accounted for how the importance of acoustic features shifts depending on genre and release year — two dimensions that capture the structural evolution of the streaming industry itself.

These questions require both theoretical and applied exploration. From a statistical viewpoint, popularity prediction can be well defined using regression, where acoustic features are tested alongside continuous covariates such as release year and categorical variables such as genre. Understanding how these relationships evolve over time is particularly important: a feature like energy may have driven popularity in 2005 but become less relevant — or even negatively associated — by 2018 as listener preferences and genre dominance shifted.

Current work in this area has largely focused on acoustic features in isolation, without accounting for the joint role of genre and release year. Li (2024) uses OLS regression using audio features with a sample of  $n = 2000$  tracks and yields an  $R^2$  value of 0.282, without including genre as a predictor. Natev (2024) quantifies hits from non-hits using neural networks from a dataset of 1,299 tracks, with accuracy of roughly 60%. Jung & Mayer (2024) have the closest methodology, employing OLS, MARS, Random Forest, and XGBoost to 30,000 tracks and account for genre as a fixed effect, finding that it is the dominant predictive factor. However, their analysis does not examine whether audio feature effects differ across genre, leaving the relationship between genre and popularity underspecified.

This paper aims to address that gap using a filtered dataset of  $n \approx 114,000$  Spotify tracks released between 2000 and 2022. This study estimates iterations of OLS regression models using the form:

$$P_i = \mathbf{x}_i^\top \beta + \varepsilon_i$$

where  $P_i$  denotes Spotify popularity score for track  $i$  and  $x_i$  represents a vector of audio

features, genre, and release year, comparing models by  $R^2$ . Genre accounts for the largest improvement in  $R^2$ , with a jump from 0.042 to 0.434. To expand on this, genre is considered alongside release year, further improving  $R^2$  to 0.584, motivating exploration of the genre-year interaction. A partial F-test confirms the significance of this ( $F(80, 318,243) = 653.72$ ,  $p < .001$ ), and the interactive model achieves an  $R^2 = 0.641$ .

The significant genre-year interaction provides motivation for deeper exploration into how release year moderates popularity across genres. The results indicate that genre is a structurally important variable with implications for statistical modeling of virality and data-driven decision making in the streaming industry. Finally, rolling coefficient analysis is applied to examine how the effects of danceability, energy, and loudness on popularity have changed over time, revealing meaningful instability in the predictive importance of acoustic features across the streaming era.

## 2 Data Overview

### 2.1 Summary of Dataset

The Spotify dataset contains data from over 400,000 Spotify tracks starting in 2000 ranging to 2023. There are over 40,000 artists spanning 81 genres. The dataset includes 19 variables per track, relating to audio features and other metadata, allowing for a full analysis.

Table 1: Dataset Summary

total_tracks	total_artists	total_genres	year_min	year_max
413495	49304	81	2000	2023

### 2.2 Popularity by Year

Table 2: Popularity by Year

year	tracks	avg_popularity
2023	15482	19.76
2022	21065	31.33
2021	21188	28.29
2020	21702	26.67
2019	21954	25.93

Table 2 reports track counts and average popularity by year in descending order. The most recent years consistently have higher average popularity scores, reflecting Spotify’s algorithmic weight on recent songs.<sup>1</sup>

Table 3: Top 10 Years by Average Popularity

year	tracks	avg_popularity
2022	21065	31.33
2021	21188	28.29
2020	21702	26.67
2019	21954	25.93
2018	22550	24.24
2017	22072	23.25
2016	15668	22.57
2015	20519	20.29
2023	15482	19.76
2014	21010	18.90

<sup>1</sup>Note the songs from 2023 tend to have a lower popularity score due to lack of time to accumulate streams.

## 2.3 Summary of Audio Features by Genre

Table 4 compares audio features in the top and bottom 5 popular genres. Higher popularity genres tend to have higher danceability and energy. Substantial variation in audio features across genre suggests that genre strongly contributes to the relationship between audio features and popularity.

Table 4: Audio Features by Genre (Top and Bottom 5)

genre	avg_danceability	avg_energy	avg_tempo	avg_popularity
pop	0.59	0.62	122.19	55.99
hip-hop	0.72	0.67	120.09	46.56
rock	0.56	0.66	124.04	46.14
dance	0.66	0.72	121.35	43.00
metal	0.43	0.88	130.66	39.60
chicago-house	0.77	0.75	125.06	3.34
detroit-techno	0.73	0.74	127.63	2.94
grindcore	0.27	0.92	118.33	2.58
tango	0.55	0.33	116.52	2.50
romance	0.43	0.30	106.92	0.58

## 2.4 Popularity by Decade

Tracks released in 2011-2020 and 2021-2023 score much higher average popularity than those in the 2000-2010 group, consistent with the patterns reported in Table 2.

Table 5: Popularity by Decade

decade	tracks	avg_popularity
2021-2023	57735	27.11
2011-2020	187855	21.81

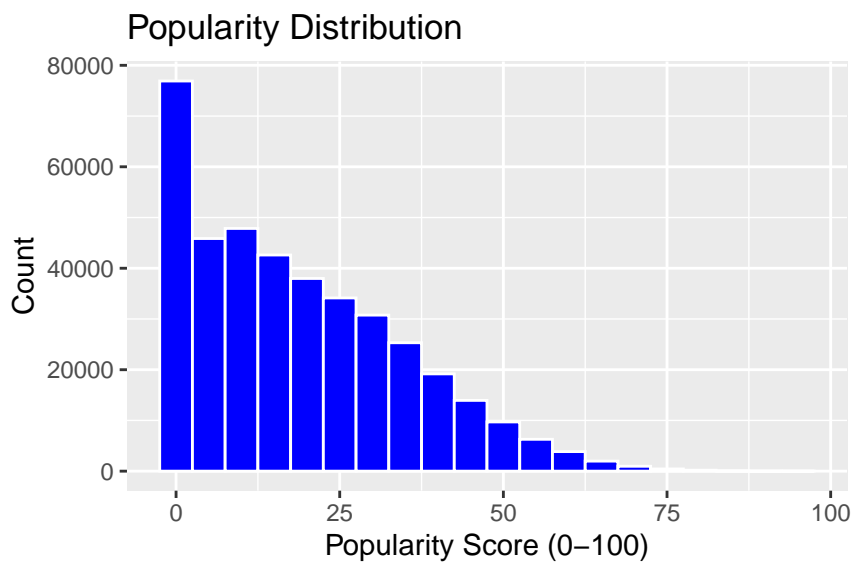
---

decade	tracks	avg_popularity
2000-2010	167905	12.71

---


## 3 Exploratory Data Analysis

### 3.1 Popularity Distribution



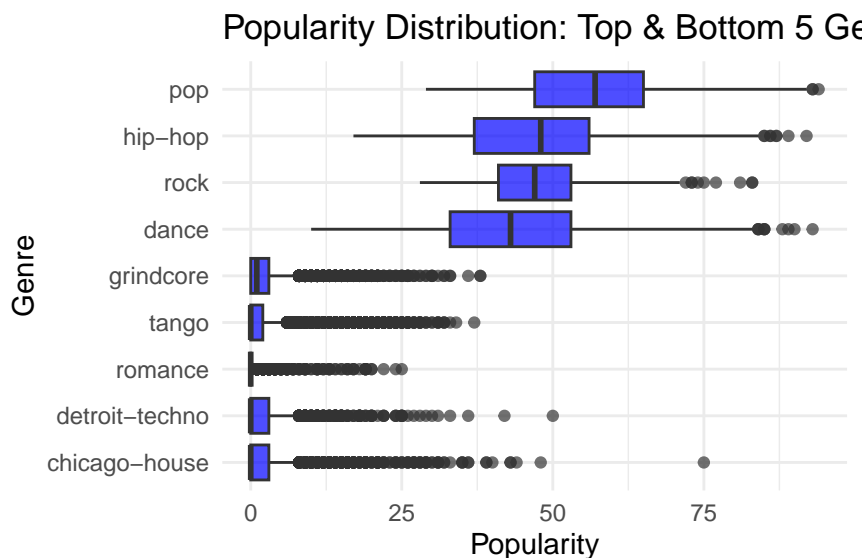
The popularity score distribution is right-skewed, with a majority of tracks falling between 0 and 20. The distribution is consistent with a long-tail structure where only a small percentage of songs accumulate the majority of streams. Songs with a popularity score of 0 are shown in the data overview for thoroughness, but removed from the sample used in modeling.

### 3.2 Popularity by Year



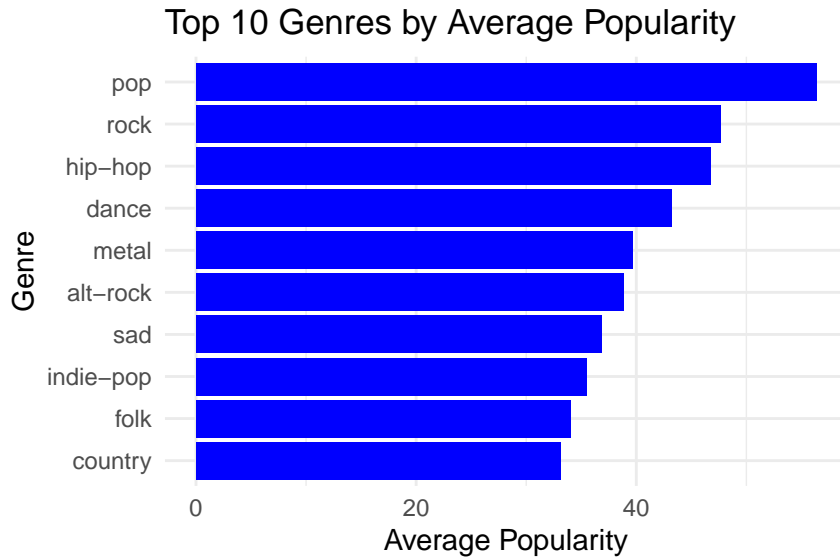
Average popularity increases monotonically from 2000 to the early 2020s, with a sharp increase after 2015. This trend aligns with Spotify’s weighting toward modern tracks and motivates the inclusion of year as a predictive factor for popularity.

### 3.3 Popularity by Genre



High popularity genres have tighter distributions, while lower popularity genres have a wider spread. The difference in central tendency and variation across genres motivates genre as a

predictive factor of popularity in the regression models.



The bar chart corroborates that genre differences in popularity are significant, with top genres having roughly double the popularity of mid-tier genres.

## 4 Model

Four iterations of OLS regression models are used, each building on the previous by incorporating additional prediction variables and the final model including interaction. All models use a training set using 80% of the sample ( $n_{train} \approx 320,000$ ), with a separate testing dataset used for testing ( $n_{test} \approx 80,000$ ). The train/test is fixed using `set.seed(123)` for reproducibility. Performance is measured using root mean squared error (RMSE) and a testing sample  $R^2$ , defined as

$$R^2 = 1 - \frac{\sum_{i=1}^n (P_i - \hat{P}_i)^2}{\sum_{i=1}^n (P_i - \bar{P})^2}$$

, where  $\hat{P}_i$  denotes predicted popularity for the  $i$  track and  $\bar{P}$  is the testing set mean.

### 4.1 Baseline Model (Audio Features Only)

Table 6: Baseline Model Coefficients (Audio Features Only)

term	estimate	std.error	statistic	p.value
(Intercept)	28.6704	0.2037	140.7655	<0.0001
danceability	10.9755	0.1808	60.7056	<0.0001
energy	-9.5863	0.1689	-56.7528	<0.0001
loudness	0.6488	0.0081	80.3121	<0.0001
liveness	-2.8795	0.1430	-20.1313	<0.0001
valence	-6.9031	0.1241	-55.6231	<0.0001

The initial baseline model predicts popularity solely on acoustic features: danceability, energy, loudness, liveness, and valence. This serves as a baseline for comparison for further testing. The  $R^2 = 0.0423$  and  $RMSE = 15.73$ , implying that acoustical variables only explain 4.2% of the variation in popularity. Loudness is the strongest predictor, which reflects the prior results (Jung & Mayer, 2024; Pareek et al., 2022). The remaining variables in the model (energy, liveness, valence) have minimal contribution. A low  $R^2$  value is expected due to acoustic properties being genre independent.

## 4.2 Improved Model (Audio Features + Genre)

Table 7: Improved Model Coefficients (Audio Features + Genre)

term	estimate	std.error	statistic	p.value
(Intercept)	21.2188	0.2245	94.4952	<0.0001
danceability	8.9097	0.1688	52.7931	<0.0001
energy	-2.5220	0.1556	-16.2123	<0.0001
loudness	0.3849	0.0073	52.5226	<0.0001
liveness	-0.8454	0.1168	-7.2407	<0.0001
valence	-6.9676	0.1086	-64.1693	<0.0001

Adding genre as a set of indicator variables produces the largest improvement in  $R^2$ , jumping from 0.0423 to 0.4340, and a decrease in RMSE from 15.73 to 12.10. This significant jump is consistent with Jung & Mayer (2024), who find that genre improves the OLS model significantly more than nonlinear additions. Pareek et al. (2022) exclude genre from their models, arguing that it does not affect prediction. The models in this paper include genre, due to the vast differences in average popularity. Genre captures major differences in popularity across musical styles, which cannot be observed solely by acoustical features due to the variation across genres.

### 4.3 Full Model (Audio Features + Genre + Year)

Table 8: Full Model Coefficients (Audio Features + Genre + Year)

term	estimate	std.error	statistic	p.value
(Intercept)	-1813.9311	5.5173	-328.7740	<0.0001
danceability	4.6683	0.1459	31.9928	<0.0001
energy	-2.2660	0.1340	-16.9114	<0.0001
loudness	0.2519	0.0063	39.8303	<0.0001
liveness	-0.2568	0.1006	-2.5537	0.0107
valence	-2.0561	0.0947	-21.7170	<0.0001
year	0.9117	0.0027	332.8244	<0.0001

The next model builds on the already improved, by considering release year, the  $R^2$  improves to 0.5807 and RMSE decreases to 10.41. The coefficient for year is positive and significant, reflecting a recency bias in Spotify’s popularity scoring, where recent songs accumulate streams faster and have more promotion in Spotify’s algorithm.

## 4.4 Interaction Model (Genre x Year)

The genre-year model assumes the popularity trend over time is constant. To compensate for this assumption, genre-year interaction terms are introduced, in the form:

$$P_i = \mathbf{x}_i^\top \beta + \sum_g \gamma_g \cdot \mathbf{1}[\text{genre}_i = g] + \delta \cdot \text{year}_i + \sum_g \theta_g \cdot \mathbf{1}[\text{genre}_i = g] \cdot \text{year}_i + \varepsilon_i$$

where  $\gamma_g$  captures genre effects,  $\delta$  is the year trend, and  $\theta_g$  lets the slope of popularity over time to differ across genres. This allows for each genre to have its own relationship with time. A partial F-test against the full model confirms that the interaction between genre and year is significant (  $F(80,318243)=653.72, p < .001$ ), meaning the popularity trends are genre specific. The  $R^2$  again increased, jumping to 0.6407 with a RMSE decrease to 9.64, the strongest performance thus far. RMSE = 9.64 means predictions deviate from the true popularity by 9.6 points on average, reflecting unmeasured variables such as algorithmic playlist tendencies and artist prominence/history.

All analyses were conducted in R (R Core Team, 2025) using the ranger (Wright & Ziegler, 2017), ggplot2 (Wickham, 2016), broom (Robinson et al., 2025), and knitr (Xie, 2025) packages. Data querying and aggregation for exploratory analysis were performed using SQL and RSQLite (Wickham et al., 2023).

## 4.5 Model Comparison

Table 9: OLS Model Comparison

	RMSE	R2
Baseline	15.73	0.0423
Improved	12.10	0.4340
Full	10.41	0.5807
Interactive	9.64	0.6407

Table 9 summarizes the sample performance across the four models. The progression from the baseline model to the interactive shows a clear trend: audio features alone provide weak predictions, genre fixed effects have the greatest prediction improvement, year contributes a meaningful increase, and genre-year interaction furthers the prediction power. Each iteration is preferred over the previous, with a higher  $R^2$  and lower RMSE. The difference between the baseline and interactive model ( $\Delta R^2 = 0.598$ ) shows the inadequacy of acoustical-only models and motivates further analysis in Section 6.

## 4.6 Residual Analysis




Figure 1: Residuals vs Fitted (Interaction Model)

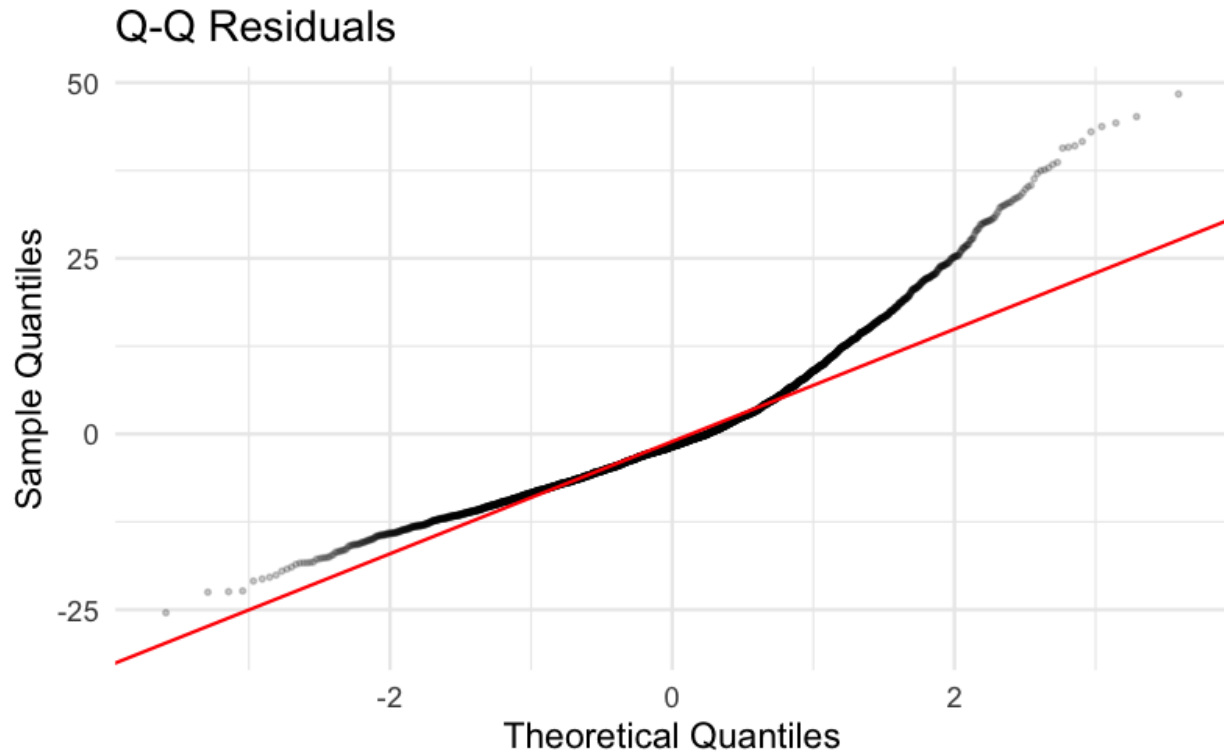


Figure 2: Q-Q Residuals (Interaction Model)

Residual analysis is explored for the interaction model, as it is the best performing OLS.

The first noticeable pattern is in the residuals versus fitted plot, with negative values for a subset of tracks, which is a limitation of applying OLS, where  $P_i \in [0, 100]$ . Songs with a low genre-year base popularity may have predicted scores below zero, which aren't meaningful. The second pattern is a wedge-shape where residuals have larger variance at low fitted values, indicating heteroskedasticity. Select observations (2064, 55290, 299316) have large residuals, meaning those tracks outperformed the genre-year predictions. The smoothed conditional expectation of the residuals is close to zero, suggest linearity.

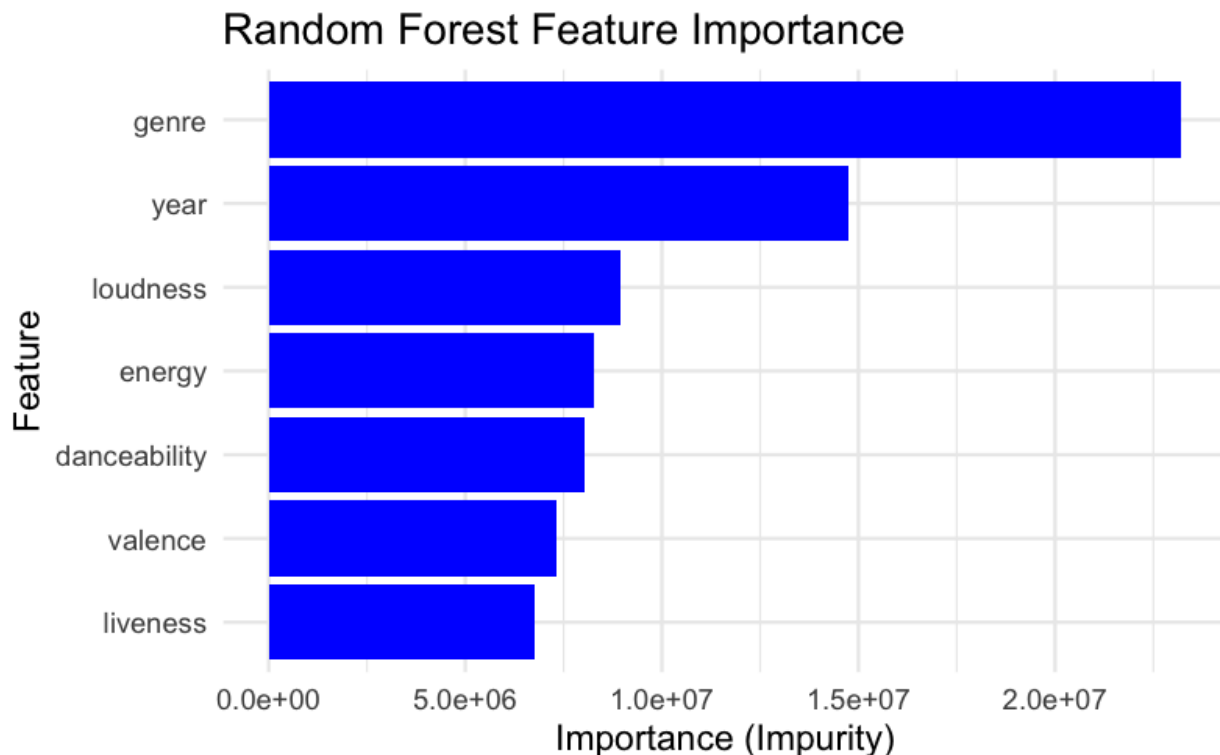
All violations (heteroskedasticity, non-normality, negative fitted values, and high leverage points) are accepted as limitations of using the OLS framework. With  $n > 300,000$ , the OLS estimator remains consistent and coefficient estimates are not invalidated. The findings from the residual plots indicates standard error and hypothesis tests should be approached with caution.  $R^2$  and RMSE are unaffected by these violations, as they measure forecasting

accuracy.

This analysis showcases the limitations of OLS. Heteroskedasticity and non-normality arise due to the bounded nature of the popularity metric ( $[0,100]$ ). Random forest regression can account for these limitations, due to its nonparametric method, there are no distribution assumptions for residuals, meaning it naturally handles boundary issues and can account for nonlinear interaction without explicit definition (Breiman, 2001). The following section uses random forest for the same prediction scheme, using the same train/test split for comparison with the OLS results.

## 5 Machine Learning Analysis

### 5.1 Random Forest - audio features + year

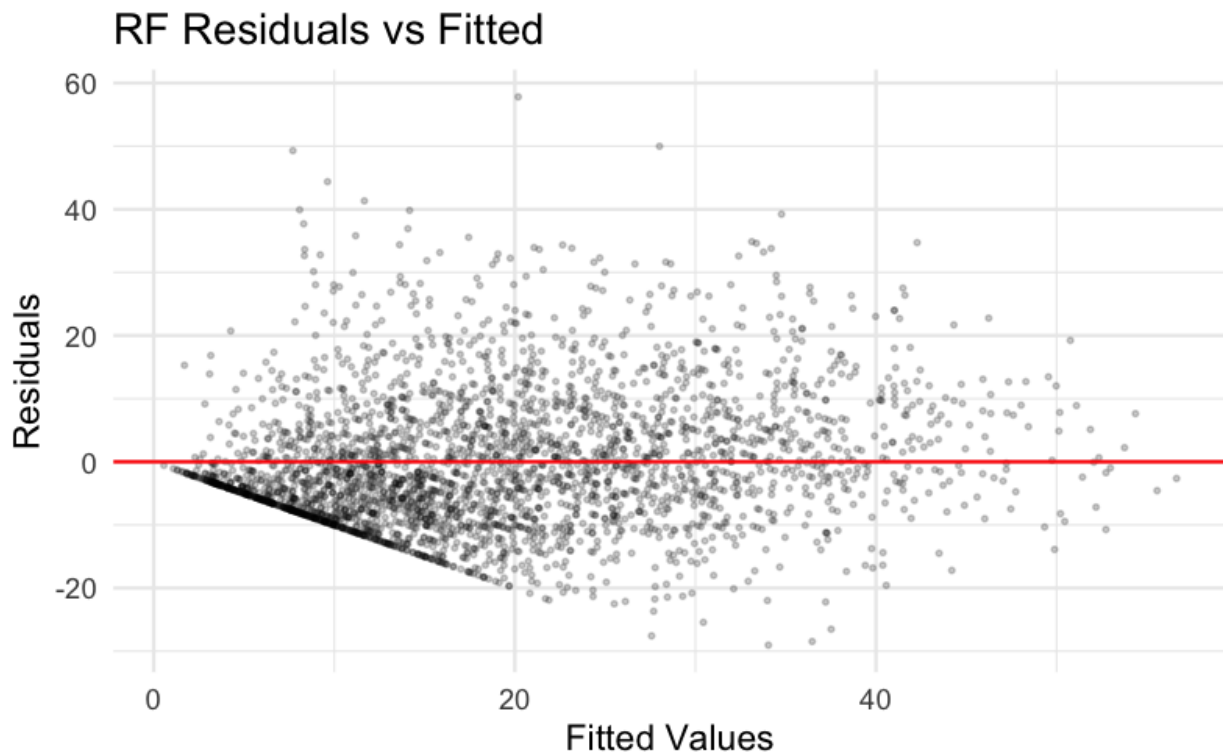



Feature importance is analyzed using impurity importance values, which measure the total node impurity reduction for each variable across all trees (Breiman, 2001). Genre is the most

important variable by a large margin, with an importance score of roughly 1.5 times year and approximately 3 times any audio feature. Year follows genre, which is consistent with the findings in the OLS models. The dominance of genre as a predictor in both the random forest and OLS models provides concrete evidence that contextual properties drive popularity prediction more than acoustical features.

Looking at the feature importance, it raises an interesting question: how do audio characteristic coefficients drive popularity over time and genre?

## 5.2 Random Forest Residual Analysis





Residual analysis for the random forest model shows improvement over the genre-year OLS model. The residual versus fitted plot no longer produces negative values, which was expected. However, heteroskedasticity still remains. The residual distribution is roughly symmetric and centered near 0, suggesting improvement over the OLS model. The right tail ends at +60, meaning the model underpredicts songs at high popularity.

### 5.3 Model Comparison

Table 10: Full Model Comparison: OLS and Random Forest

	RMSE	R <sup>2</sup>
OLS Baseline (Audio Only)	15.73	0.0423
OLS Improved (Audio + Genre)	12.10	0.4340
OLS Full (Audio + Genre + Year)	10.41	0.5807
OLS Interaction (Genre $\times$ Year)	9.64	0.6407

Random Forest (Audio + Genre + Year)	10.41	0.5952
--------------------------------------	-------	--------


Table 10 shows sample performance for all five models.

The random forest model performs between the OLS full and OLS interaction model. This result is unexpected, due to the nonparametric nature of the random forest. However, the residual analysis corroborates earlier findings that genre and year dominate audio characteristics in prediction.

## 6 Audio Feature Effects on Popularity Over Time

The significant genre-year interaction ( $F(80, 318,243) = 653.72, p < .001$ ) implies that audio feature effects on popularity are not constant across time or genre. Rolling coefficient analysis formalizes this by estimating how the marginal effect of top audio feature shifts by year.

### Effect of Audio Features on Popularity Over Time



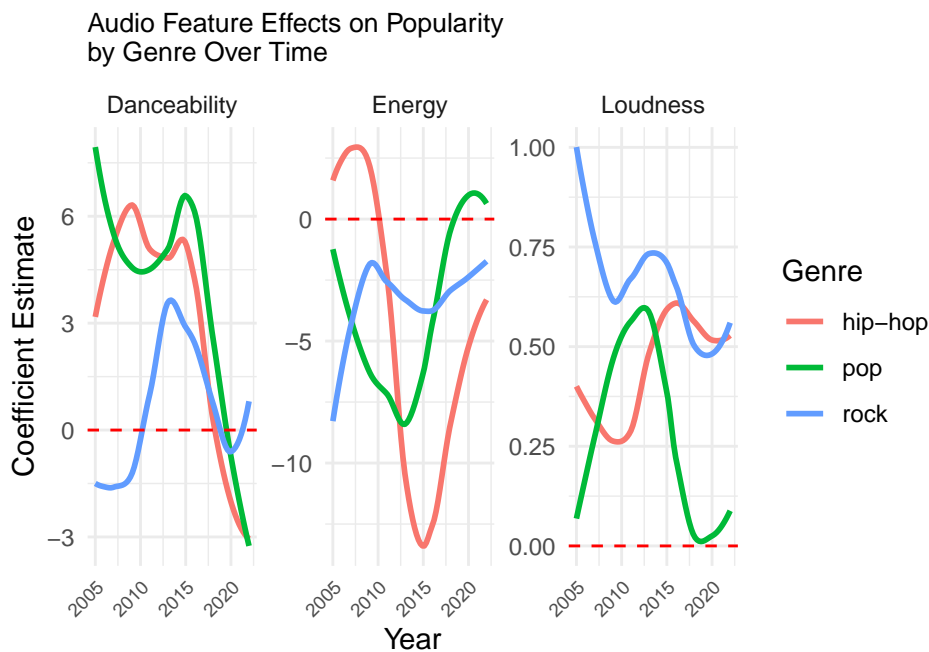
Danceability (red) has consistently had a positive effect on popularity, showing a slight dip around 2010-2012, and peaking around 2016 with the rise of streaming and genres such as electronic dance music. It shows a slight decline in recent years but still remains positive. Energy (green) started around 0 and started to dip into negative territory, with a low around -3.5, reflecting a shift in popularity to more slower-paced songs. It has started to resurface,

and requires further investigation with modern data. Finally, loudness (blue) peaked in 2011-2012, likely due to the “loudness wars” and started to decrease. However, it has not passed 0, indicating an overall positive in popularity contribution.

Since the genre-year interaction OLS model had the highest  $R^2$  and lowest RMSE, it is natural to explore audio characteristic effects on popularity based on genre and year, looking at the top three genres for average popularity.

Table 11: Top 3 Genres by Average Popularity

Genre	Avg. Popularity	n
pop	56.85	424
rock	47.11	250
hip-hop	46.54	1124



The three top genres (pop, rock, hip-hop) all show a positive danceability effect on popularity, but have very different trajectories. Pop again peaked around 2015 with a sharp decline after, suggesting that danceability no longer has the same contribution to popularity in pop music. Hip-hop has remained consistently high, which is expected for the genre. Rock’s danceability effect started negative, turned positive (2010-2012) and currently sits around 0.

There has been a dramatic decline in energy’s popularity effect after 2010. Pop, rock, and hip-hop have all had a significant decline, with hip-hop being the most severe. There was a slight resurgence around 2018, which has again declined.

Loudness’s contribution to popularity peaked around 2005 and started to decline after 2010. Pop has been relatively stable and hip-hop has had a slight increase, consistent with the growing genre.

Taken together, these trajectories provide a strong interpretation of the genre-year interaction model in Section 4. The significant partial F-test ( $F(80, 318,243) = 653.2, p < .001$ ) reflects not only that genres differ in popularity, but also that the acoustic factors within each genre have evolved with the genre. Models that utilize a common year trend across varying genres, misspecifies the data generating process, which is what the baseline and improved OLS models do. The rolling coefficient analysis provides the interpretive foundation for why the interaction model outperforms the others.

## 7 Results

The OLS iterative models find a significant jump in prediction power with the addition of genre, year, and interaction between genre and year, with the greatest difference being  $\Delta R^2 = 0.598$ . Audio features alone aren’t strong enough to predict popularity because of the vast distribution of genres and different styles dominating different time periods being clustered into a non-differentiated set. The largest increase was with the sole addition of genre ( $\Delta R^2 = +0.3917$ ), with the next largest gain being year ( $\Delta R^2 = +0.1467$ ). The interaction between genre and year is significant ( $F(80,318,243)=653.72, p<.001$ )).

The genre-year interaction model is the strongest with  $R^2 = 0.6407$  and the random forest is the second strongest with  $R^2 = 0.5952$ . Residual analysis from the random forest model indicates predictor strength: genre > year > loudness > energy > danceability > valence > liveness. Despite the under performance of the random forest model, the findings are consistent across models: genre and year are the strongest predictors of popularity, and the

interaction between genre and year allows the best prediction due to different genres being prevalent at different period of time.

The strongest model (genre-year interaction) has a RMSE of 9.64 on a 0-100 scale, meaning the predictions deviate from the true popularity index by 9.64 points on average. The unexplained variance reflects factors not captured from the dataset such as algorithmic boosts in playlists, artist prominence, and random variation. Audio features and genre-year attributes are significant, but popularity is ultimately shaped by external factors not captured by the dataset.

Rolling coefficient analysis reveals that audio feature effects on popularity differ over time and over genre. The energy coefficient declines significantly across hip-hop, rock, and pop between 2013 and 2017, reflecting a structural shift towards lower energy content on music streaming platforms. Danceability effects on popularity differed greatly between genres, with hip-hop maintaining a consistently positive contribution and pop turning to a negative contribution in recent years. Loudness effects were highest in rock around 2005 but declined after 2010, which is expected due to rock's overall decrease in streaming. These genre specific patterns provide an interpretation for the genre-year interaction seen in the OLS model.

## 8 Conclusion

This study shows the dominance of genre as the strongest predictor of a song's popularity on Spotify, contributing the largest gain in variation across all models. Genre and year have a strong interaction, indicating that genres have a different trajectory over time. Some genres have become more popular recently, while others have declined. Random forest corroborated these findings. For professionals in the music industry and artists, these findings suggest that strategic release timing can significantly contribute to a song's success. Artists should pick the acoustical features of their choosing, but carefully consider when to release said song and how they frame/market it. Despite the positive increase with genre and year interaction, the best model still leaves 35.9% of variance unexplained.

Coefficient analysis further reveals that audio features' effect on popularity are not fixed over time or genre. Energy's effect on popularity after 2010 became negative for hip-hop, rock, and pop (the three genres with the highest average popularity). This is likely due to a structural shift towards playlist curation that favors low energy content. Danceability and loudness have their own genre specific trajectories, reflecting how the features that dominated popular songs at one time, may be hurting popularity currently. For artists and labels, these findings suggest that acoustical and stylistic features should be carefully considered and not solely chosen based on norms of the past.

OLS and random forest models have several limitations within this dataset. The popularity index is bounded ( $[0, 100]$ ) which results in heteroskedasticity and produces negative values. Additionally, the models identify correlations of popularity, not causation of popularity. The nature of Spotify's popularity score has a recency bias, so older songs are disadvantaged and underrepresented compared to newer tracks. Available data does not account for playlist and suggestion algorithms.

Future research could replicate these models but for artist specific datasets, as well as using the Spotify API to pull the most recent data. Additionally, taking into account social media traction for artists and playlist inclusion could test to see if the models improve. It is also worth exploring whether similar streaming platforms such as Apple Music have similar results.

## References

- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- Jung, N. S., & Mayer, F. (2024). *Beyond beats: A recipe to song popularity? A machine learning approach*. <https://arxiv.org/abs/2403.12079>
- Li, K. (2024). Predicting song popularity in the digital age through Spotify's data. *Theoretical and Natural Science*, 39, 68–75. <https://doi.org/10.54254/2753-8818/39/20240600>

- Natev, A. (2024). Hit song prediction through machine learning and Spotify data. *Proceedings of the 10th Student Computing Research Symposium (SCORES'24)*, 57–60. <https://doi.org/10.18690/um.feri.6.2024.13>
- Pareek, P., Shankar, P., Pathak, P., & Sakariya, N. (2022). Predicting music popularity using machine learning algorithm and music metrics available in Spotify. *Journal of Development Economics and Management Research Studies*, 9(11), 10–19.
- R Core Team. (2025). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. <https://www.R-project.org/>
- Robinson, D., Hayes, A., Couch, S., & Hvitfeldt, E. (2025). *Broom: Convert statistical objects into tidy tibbles*. <https://doi.org/10.32614/CRAN.package.broom>
- Wickham, H. (2016). *ggplot2: Elegant graphics for data analysis*. Springer-Verlag New York. <https://ggplot2.tidyverse.org>
- Wickham, H., Müller, K., James, D., & Falcon, S. (2023). *RSQLite: SQLite interface for r*. <https://CRAN.R-project.org/package=RSQLite>
- Wright, M. N., & Ziegler, A. (2017). ranger: A fast implementation of random forests for high dimensional data in C++ and R. *Journal of Statistical Software*, 77(1), 1–17. <https://doi.org/10.18637/jss.v077.i01>
- Xie, Y. (2025). *Knitr: A general-purpose package for dynamic report generation in R*. <https://yihui.org/knitr/>