

THE FAV CORPUS: AN AUDIO DATASET OF FAVORITE PIECES AND EXCERPTS, WITH FORMAL ANALYSES AND MUSIC THEORY DESCRIPTORS

Ethan Lustig

Ethan@EthanLustig.com

David Temperley

Eastman School of Music, University of Rochester

ABSTRACT

We introduce a novel audio corpus, the FAV Corpus, of over 400 favorite musical excerpts and pieces, formal analyses, and free-response comments. In a survey, 140 American university students (mostly music majors) were asked to provide three of their favorite 15-second musical excerpts, from any genre or time period. For each selection, respondents were asked: “Why do you love the excerpt? Try to be as specific and detailed as possible (music theory terms are encouraged but not required).” Classical selections were dominated by a very small number of composers, while the pop and jazz artists were diverse. A thematic coding of the respondents’ comments found that the most common themes were melody (34.2% of comments), harmony (27.2%), and sonic factors: texture (27.6%), instrumentation (24.3%), and timbre (12.5%). (Rhythm (19.5%) and meter (4.6%) were less present in the comments.) The comments cite simplicity three times more than complexity, and energy gain 14 times more than energy decrease, suggesting that people’s favorite excerpts involve simple moments of energy gain or “build-up”. The complete FAV Corpus is publicly available online at EthanLustig.com/FavCorpus. We will discuss future possibilities for the corpus, including potential directions in the spaces of machine learning and music recommendation.

1. INTRODUCTION

Why do we like the music we like? Perusal of the range of options in a record store or streaming platform, or of live music offerings in a large city, shows the enormous diversity of musical taste among individuals. Research has probed some of the factors involved in this variability, such as gender, age, personality, social identity, cultural background, and musical training [1-6]. Still, there seems to be general agreement that particular pieces of music are especially “good.” Certain hymns, Christmas carols, folk songs, and classical pieces remain favorites across decades and centuries; certain popular songs cause world-wide and lasting explosions of enthusiasm. It seems, also, that specific sections or moments within these pieces are especially pleasurable, giving rise to what are sometimes called

peak experiences [7-8]. Our own personal reflections certainly confirm this, and anecdotally, there seems to be at least some agreement as to what the “best” moments of a piece are. But what makes a certain part of a piece especially enjoyable?

Music psychology has begun to address this issue, though in tentative and exploratory ways. Most of this research has focused on the physiological manifestations of peak experiences, such as chills, which have been shown to correlate with pleasure. A pioneering study by Sloboda [9] asked participants to identify passages causing strong physiological effects—what he called “thrills” (p. 110, after [10])—and to describe the nature of those responses. More recent studies follow Sloboda’s model in having participants identify pieces that cause physiological responses, especially chills, and then probing the possible causes and correlates of these responses: neurological correlates [11, 12], musical elicitors [13, 14], and self-reported perceptual correlates such as the perceived sadness or happiness of the music [15]. With regard to musical elicitors of chills, studies have found many factors including sequences, appoggiaturas, new or unexpected harmonies, crescendi, climaxes, sudden dynamic or textural changes, and entrances of instruments [9, 13-15]. Also deserving mention is a large project by Gabrielsson and Wik [16] focusing on the effects (physical, emotional, and cognitive) of “strong experiences” of music (p. 158). Musical elicitors are mentioned only briefly and in very general terms: “instruments, rhythm, melody, harmony, musical form, performance qualities etc.” (p. 198; see also [17], p. 568).

In this study we offer a novel approach to the study of peak experiences and the musical factors that elicit them. In contrast to the exploratory research cited above, our study takes a systematic, survey-based approach. Our conception of peak experience is close to Maslow’s [7]—a primarily internal albeit not physiological, intensely positive experience—and falls within the fairly broad range of ways that the term is used [8].

Our project differs from most research on peak experiences by focusing on passages of music directly reported to be strongly liked, rather than those causing chills and other physiological responses. While chills have generally been shown to coincide with pleasurable experiences [12, 13], they may not always do so; conversely, one can certainly get great enjoyment from a musical passage without experiencing chills.

In a survey, 140 respondents identified favorite musical excerpts. Respondents also provided free-response comments explaining their choices, and we provide a content



analysis of these comments. We also present a publicly available corpus, the FAV Corpus, which includes audio files of excerpts and complete pieces, formal analyses of a subset of the pieces, and the respondents' free-response comments.

2. METHOD

2.1 Participants

In 2017, 140 students at the University of Rochester (New York) were given a survey regarding their favorite musical excerpts. Approximately 85% of the respondents to the survey were students at the Eastman School of Music (a division of the university) and were therefore music majors. The remaining 15% were students in an introductory music psychology course; while students in this course were mostly not music majors, the course required basic knowledge of music theory as a prerequisite. While students with music-theory training may not be representative of the broader population, we deliberately chose them for their ability to articulate the musical reasons for their preferences, with regard to matters such as harmony, rhythm, and form. On average, the respondents had 11.1 years of music training on a musical instrument (including voice) ($SD = 4.2$). There were 73 females, 63 males, and four who preferred not to say. The average age of the respondents was 19.7 years old, with a range of 17 to 29 years old ($SD = 1.9$). Respondents received extra credit points in their courses for participation. The survey received ethical approval by the Institutional Review Board of the University of Rochester.

2.2 Data Collection

The survey asked each respondent to identify “three of your favorite excerpts of music... in any style and from any time period.” For each excerpt, they were instructed to provide a URL (web address) to a recording of the piece/song/movement on YouTube or Spotify. We used the phrase “piece/song/movement” to avoid stylistic bias, but hereafter we will refer only to “pieces.” Respondents were then asked to “identify the 15-second excerpt that’s your favorite” by providing start- and end-points for the excerpt in relation to the recording. The choice of 15 seconds was fairly arbitrary. We chose it, in part, because it roughly corresponds to the length of some of our favorite musical passages.

Following each selection, respondents were prompted to write a response to the following question: “Why do you love the excerpt? Try to be as specific and detailed as possible (music theory terms are encouraged but not required).” We take the term *love* to indicate a high degree of liking or preference, similar in meaning to *enjoy* or *greatly like*. Our mention of “music theory terms” was aimed at encouraging respondents to identify the musical features giving rise to their preferences. There is a possible downside to this wording; by drawing attention to our own

music-theoretical background, it may have steered respondents toward pieces or excerpts that they thought were theoretically “respectable” in some way. However, the huge stylistic variety of the chosen excerpts (described below), including many from very recent popular music, suggests to us that this was not a concern for many respondents. Additionally, respondents were asked to choose between either “I enjoy this excerpt much more than the other parts of the piece” or “I enjoy this excerpt about as much as the other parts of the piece.”

2.3 Creating the Corpus

Recordings of the complete pieces provided by the respondents were extracted from the YouTube/Spotify URLs and saved as WAV audio files; audio files were also made of each preferred 15-second excerpt. In some cases, the beginning of the internet recording did not correspond to the true beginning of the piece. To adjust for this, any time before the beginning of the piece was subtracted from the timepoints of the preferred excerpt, so that the adjusted timepoints indicated the excerpt’s location in relation to the piece. In about 8% of cases, the chosen excerpt was not exactly 15 seconds long, but usually just a few seconds shorter or longer. In such cases, the excerpt was converted to a 15-second excerpt with the same midpoint as the chosen excerpt. (For example, 0:00-0:25 would be converted to 0:05-0:20.) For more detail about this process, see [18].

The corpus, which we call the FAV Corpus, is publicly available at EthanLustig.com/FavCorpus. The corpus contains 420 items (three excerpts from each of the 140 respondents). A spreadsheet indicates, for each item, (a) the respondent’s number, which had been assigned arbitrarily, (b) the excerpt number for that respondent (1, 2, or 3), (c) the artist and title of the piece, (d) the style and historical era or year (explained below), (e) the duration of the piece, (f) the timepoints of the preferred excerpt, (g) whether the respondent indicated that they enjoyed the excerpt “much more than” [A] or “about as much as” [B] the rest of the piece, and (h) the respondent’s comment about why they liked the excerpt. In what follows, we indicate excerpts by respondent and excerpt number; for example, Respondent 1’s three excerpts are 1_1, 1_2, and 1_3. We also provide sound files for both the complete pieces and the preferred 15-second excerpts.¹

3. RESULTS

3.1 Stylistic Content of the Corpus

The distribution of styles and artists in the corpus was examined. While this is not the main focus of the current study, it provides a window into the musical tastes and passions of students at an American music school in 2017 (recall that roughly 85% of respondents were music students). Each excerpt was categorized as classical (49.5%), pop (41.8%), or jazz (8.7%). For most excerpts, classification was clear; there were a few borderline cases, such as jazz-rock fusion pieces. The most popular artists in the survey

¹ Due to reasons such as invalid web links, respondent errors, etc., the actual corpus contains 399 excerpt audio files and 402 piece audio files. (See EthanLustig.com/FavCorpus for details.)

are listed in Table 1, with the number of excerpts for each. Following convention, for classical works, we identify the composer as the artist; for jazz and pop, we identify the performer(s) as the artist. Table 1 alone might give the impression that respondents strongly favored classical music, but the style statistics just cited show otherwise; the preponderance of classical composers in Table 1 indicates, rather, that classical selections were dominated by a small handful of artists, while pop and jazz selections were much more widely dispersed.

Composer	# Excerpts
Bach	17
Brahms	14
Beethoven	12
Tchaikovsky	10
Rachmaninov	9
Mahler	6
Kendrick Lamar	6
Debussy	6
Sibelius	5
Handel	5

Table 1. Artists (composers/performers) most represented in survey.

Among the classical excerpts, 12.1% were from the Baroque period (1600-1749), 8.5% Classical (1750-1819), 31.7% Romantic (1820-1899), and 47.7% 20th/21st-century (1900-present). (Each composer was assigned to a single period, based on their years of greatest activity.) Again, the large number of 20th/21st-century selections is not reflected in Table 1 since they are distributed over a much larger number of composers. We also observed that many of these 20th/21st-century composers were toward the conservative end of the stylistic spectrum; the most popular was Rachmaninoff, with nine excerpts. For the pop and jazz selections, we identified the year of release of each recording. The pop selections strongly favored recent music: 69.0% were from 2010–2017 (more than half of these from 2016–2017 alone), and 17.9% from the 2000s. Jazz selections had a weaker bias toward recent music, with 31.4% of selections from 2010 through 2017.

3.2 Formal Analysis

One of us (David Temperley) did a formal analysis of a subset of pieces in the corpus. He did not know which excerpts were preferred when doing the analysis. The subset consisted of pieces in which respondents had said that they liked their preferred excerpt “much more than” the rest of the piece; this yielded a set of 127 pieces (about 30% of the survey responses).² The recordings of the pieces were divided into sections to the nearest second, and the sections were given formal labels, as the genre warranted (for instance, P = primary theme for a sonata-form piece; V (verse) and CH (chorus) for pop songs). It was assumed that each section continued until the beginning of the next

section, so that each piece was exhaustively partitioned into sections. As an arbitrary constraint to simplify the analysis, no section was allowed to be less than 15 seconds long. Two main criteria were used for determining the location of formal sections: change and repetition. A significant change in any musical parameter, such as harmony, melody, instrumentation, texture, meter, or rhythmic pattern, was considered to make a good candidate for a section break. Repetition could also define sections: for instance, the return of the opening theme in a sonata-form movement might define a new section beginning even in the absence of obvious local changes. Repetition of the same label signified exact or slightly modified repetition; for example, V would be used for two verses of a pop song, with different lyrics and perhaps some changes in instrumentation, but mostly similar melody and harmony. For more substantially modified repetitions, numbers were used (e.g., V1 and V2 for two verses that had significantly different melody or harmony). See [18] for more detail about the annotation system.

We analyzed the preferred excerpts in relation to their location within the piece. First, we wondered if people tend to choose excerpts that are near section boundaries. For each preferred excerpt, in the set of 127 excerpts for which formal analyses were available, we found the temporal distance between the midpoint of that excerpt and the closest formal section boundary; we then performed the same process for random 15-second excerpts from the same pieces, repeating the process 10 times to mitigate the effect of extreme values. (One piece had a 7-minute section that seemed to create outliers in the data; this piece was removed from the analysis.) Midpoints of preferred excerpts have an average (absolute) temporal distance from the nearest section boundary of 11.41 seconds, while for midpoints of random excerpts, the distance is 13.67 seconds—a modest but significant difference ($t(168.73) = -2.46, p < 0.01$). Thus, preferred excerpts show a slight tendency to be located near formal boundaries. A total of 49.2% of the preferred excerpts actually contain a section boundary; among the random excerpts, only 37.6% do. We then re-analyzed the same distances as signed values, to see whether preferred excerpts tend to be near the beginning or end of a formal section. For preferred excerpts, the mean signed difference between the midpoint and the nearest boundary is 2.90 (i.e., on average, the midpoint occurs 2.90 seconds after the boundary), significantly greater than zero (one-sample t -test, $t(125) = 2.22, p < 0.05$). This indicates a slight tendency to choose excerpts near the beginning of a section rather than near the end, or, perhaps, overlapping more with the beginning of a section than with the end of the previous one.

Finally, we examined the location of each excerpt in relation to the piece as a whole. For this analysis, we used all 399 excerpts in the corpus. Each excerpt received a value for its proportional position in the piece, where 0 would be at the very beginning, and 1 would be at the very end. The

² Altogether there were 137 eligible pieces, 10 proved impossible to analyze into formal sections, because there was no large-scale repetition and no clear moments of change demarcating reasonably-sized sections. Some of these were contemporary pieces; others were Baroque pieces,

for example imitative textures with a rapid or seamless alternation between subject entries and episodes.

mean value was 0.46; clearly there was not a strong bias toward choosing excerpts early or late in the piece.

3.3 Content Analysis of Comments

As mentioned earlier, respondents were asked to comment on their reasons for liking each excerpt in their own words. Responses varied from a few words to several sentences. While a few responses were flippant or minimal, a great many respondents showed enthusiasm for the task and took considerable effort in explaining their choices. We did a content analysis of the respondents' comments. One of us (David Temperley) coded all 420 comments, identifying 17 themes that seemed to appear repeatedly in the comments. We then provided a list of the 17 themes and their definitions (Table 2) to an independent coder (a music theory Ph.D. student at the Eastman School of Music) and asked him to assign themes to the comments using that list. Each comment could be tagged with any number of themes (including zero). In choosing themes, both coders aimed to represent the respondents' actual reasons for liking the excerpts, as opposed to aspects mentioned simply to aid reference, although this distinction was not always easy to make. For example, a comment like "I love the violin melody" was encoded as MEL (melody) rather than INS (instrumentation).

BIO Autobiographical connection: references to the respondent's past experience with the piece or excerpt, OR incidents in their life that it reminds them of for any reason.

COM (+/-) Complexity (or its opposite, simplicity).

DYN (+/-) Dynamics.

EN (+/-) Energy. Energy level in music is thought to be conveyed by such as dynamics, register, rhythmic activity, and textural thickness; an increase in any of these dimensions could create a rise in energy. However, when the change is described in these more specific terms (e.g. dynamics) it can be coded in that way; EN should be reserved for more general descriptions of energy change or level, e.g. "buildup" or "climax".

HAR Harmony: includes harmonic progression, function, or chord quality; also tonality (e.g. modulation), mode (major/minor), and dissonance/consonance.

INS Instrumentation: choices of instrument or instrument combinations; also includes general uses of an instrument (e.g. "I like the clarinet in a high register"), or special timbral effects prescribed by composer, e.g. extended techniques; also synthesized parts in popular music textures. (Compare to TIM).

INT Interpretation (e.g. expressive timing; also general statements about beauty/expressiveness of a performance or quality of performer).

LJR Lyrics.

MEL Melody: the main melody in this particular part of the piece. Also includes improvised solos, e.g. in jazz.

MET Meter (incl. tempo).

PHY Mentions of a physical or physiological response to the music.

RET Return of earlier thematic material.

RH Rhythm. Includes references to general rhythmic feel,

e.g. "groove".

SUR Explicit mentions of surprise or denial of expectation.

TEX Texture: a catch-all category including aspects of pitch-rhythmic patterns other than melody, such as details of accompaniment or bass lines, chord voicings, or polyphonic patterns.

TIM Timbre: when credited to performer (e.g. a singer's tone), or synthesized/electronic sounds that are not a consistent part of the texture. (Compare to INS).

VIR Virtuosity (or just proficiency, i.e. playing a very difficult bit accurately; also intonation).

Table 2. Themes and definitions used in content analysis.

Agreement between the two coders regarding the assignment of each theme was measured using Cohen's kappa, where 1.0 would indicate that the two coders assigned the theme to exactly the same comments. Agreement levels varied between 0.37 and 0.84, depending on the theme, and thus were mostly in the range of *moderate* or *substantial* according to Landis and Koch's [19] rubric. In what follows we discuss the results of this content analysis. We also analyzed word frequencies in the comments, grouping together similar words such as "simple," "simpler," and "simplicity". We include some results of that analysis in the following discussion.

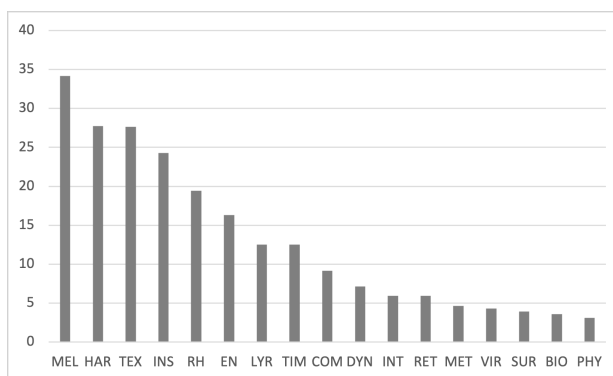


Figure 1. The percentage of respondents' comments identified with themes in the content analysis. For explanation of abbreviations, see Table 2.

Figure 1 shows the percentage of occurrences of each theme in the comments. The counts of each theme were averaged between the two coders. The frequent mentions of melody (MEL, occurring in 34.2% of comments) and harmony (HAR, 27.2%) indicate the importance of the pitch domain in respondents' preferences. Rhythmic factors—rhythm (RH, 19.5%) and meter (MET, 4.6%)—were less important, though it should be remembered that melody has a rhythmic aspect as well. Notably, the word "groove" occurred 20 times—confirming the widely held view that this is a significant factor in musical enjoyment [20, 21]. What might be called sonic factors were also mentioned frequently: texture (TEX, 27.6%), instrumentation (INS, 24.3%), and timbre (TIM, 12.5%). There were comparatively few mentions of performance aspects: interpretation (INT, 6.0%) and virtuosity (VIR, 4.3%).

Given that a large majority of respondents were majoring in classical music performance, we had expected these factors to weigh more heavily. Lyrics (LYR) were mentioned in 12.5% of comments, and autobiographical factors (BIO, connections with the respondents' life experience) in just 3.6%. The PHY theme, physiological responses (such as chills), was mentioned in just 3.1% of comments. It is possible that the survey instructions—which encouraged the use of music-theory terms—steered respondents' attention towards musical features and away from autobiographical and physiological factors.

Three of the themes—complexity (COM), energy (EN), and dynamics (DYN)—were parametric: They could be subscripted as “+” (indicating an increase or relatively high level) or “-” (a decrease or relatively low level), though this was optional. While there were 6.5 instances of COM+ in the comments, there were 22 instances of COM- (again, theme counts reported here and below are averaged across the two coders' analyses). This result suggests that respondents favored moments of relatively low or decreasing complexity. Our analysis of word frequencies also supports this view: “simple” (and related words) occurs 32 times in the comments, while “complex” (and its variants) only occurs 11 times. Related to this, the words “tension/tense” and “resolution/resolve” were used about equally often (19 and 18 times, respectively). However, seven of the comments mentioning tension refer specifically to the resolution of the tension (sometimes using other words like “relax,” “release,” or “relief”); in the remaining cases, the tension seems to be valued in itself.

The energy (EN) theme shows an even stronger parametric tilt than complexity: 35.5 of its mentions are EN+, while only 2.5 are EN-. Energy is often treated as more or less synonymous with the arousal/activation dimension in Russell's [22] two-dimensional model of emotion, and this in turn has been associated with musical parameters such as loudness, pitch register, and rhythmic activity [23].³ Note from Table 2 that this theme reflects general references to energy, as opposed to mentions of energy-invoking musical dimensions such as dynamics, rhythm, or texture. The dynamics (DYN) theme also showed a parametric tilt, marked “+” eight times and “-” only three times. Analysis of word frequencies shows further evidence of a preference for increasing energy. For example, the word “build” and related words such as “build-up” occurs 47 times. It is not obvious what the opposite of “build” would be, indicating a general *decrease* in energy level; one thinks of such words such as “decrease,” “decline,” “fade,” “wane,” “subside,” and “dwindle.” None of these words occurred even once, except “fade,” which occurred just three times.⁴ Several other frequent word categories indicate an increase or peak in energy, such as “climax/climac-

tic” (used 28 times), “power(ful)” (27 times), and “crescendo” (9 times; “diminuendo” is never used and “decreasing” just once).

In this connection, a result from our analyses of formal structures, described earlier, is relevant. In pop songs, which nearly always contain both choruses and verses, respondents' preferred excerpts were more often in choruses (13 times) than verses (7 times). (Recall that our analysis of formal structures included only about 30% of the survey responses.) Respondents' comments also mentioned choruses (44 times) much more often than verses (18 times).⁵ It has been observed that choruses tend to be higher than verses in the “energetic” dimensions mentioned above, such as pitch register and textural thickness [24, pp. 39-40]. Thus, several patterns in our data point to an increase in energy as an important elicitor of musical pleasure.

Perusal of the comments suggests other possible themes as well. For instance, many comments contain terms or phrases that could be described as emotional. In the first 20 comments, we see “aggressive” (1_2), “raw emotion” (2_3), “intensity” (2_3), “exciting” (4_2, 6_2), “[the singer] let[s] emotions loose,” (5_2) “dramatic” (6_1), and “triumphant” (6_2). In many cases, such terms are used to describe a specific aspect of the music that could also be encoded in some other way: for example, “a triumphant theme” (MEL); “the buildup is very exciting” (EN). Another issue is the distinction between induced and perceived emotion [25]. Sometimes the distinction is clear—“It is insanely happy” is perceived emotion, “[It] always makes me so happy” is induced emotion—but not always: if a passage is described as “exciting” or “relaxing,” is that induced or perceived emotion? If induced emotion is included in the “emotion” theme, one could potentially include a large number of comments implying a positive emotional reaction: for example, “I love the cellist's interpretation.” Indeed, one might say that such a reaction is implicit in all of the comments, given the nature of the task.

4. DISCUSSION

In our study, 140 college students, mostly music students, identified three of their favorite 15-second passages of music. One result emerging from our analysis of the survey comments was a preference for passages that increase in energy—often described by respondents as “builds” or “build-ups,” or as sections that “build.” As noted earlier, energy in music is generally associated with parameters such as loudness, pitch register, and rhythmic activity. It also seems intuitive to us, although this does not seem to have been widely studied, that textural thickness is also associated with energy, perhaps partly because a thickening of texture implies greater loudness, whether or not the loudness actually increases. Our finding that increases in

³ In experiments on music and emotion, manipulations in the temporal dimension usually involve changing the speed of a melody, and are therefore described (correctly) as variations in tempo (for a survey, see [23]). Within a piece, however, the tempo (i.e., the speed of the main beat) rarely changes, except for small fluctuations; temporal variation is more likely to involve changes in rhythmic values (e.g., from a quarter-note texture to a 16th-note texture). In both of these cases, though, the variations involve a change in the temporal density of events; if an increase in tempo conveys an increase in energy, it seems likely that an increase in rhythmic activity over a fixed tempo would also do so.

⁴ Some of these words, such as “build,” “decrease,” and “fade,” could be either nouns or verbs; we counted both, including all verb forms. The word “drop” is also of interest; it occurs 12 times, as noun or verb, but only five of those uses could be taken to refer to energy level. Sometimes the term is used to refer to the re-entrance of the kick drum in a pop or EDM song.

⁵ One might wonder if choruses are more frequent than verses in our corpus, and therefore take up more time. Actually they do not: choruses take up a total of 1769 seconds, in the portion of the corpus that was formally analyzed; verses take up a total of 1977 seconds.

energy are often pleasurable accords well with other work on peak musical experiences that has linked them to crescendo and increases in texture [13-15, 26]. It also appears that there is a strong preference for passages perceived as having relatively low or decreasing complexity and tension, compared to passages perceived to be high in complexity or tension. This is in line with Meyer's [27] observation that in music, "The greater the buildup of suspense, of tension, the greater the emotional release upon resolution" (p. 28) and Huron's [28] idea of "contrastive valence" (p. 39).

Earlier studies have found that a wide range of factors affect peak experiences [9, 13, 14], and this is apparent from the free-response comments in our survey. The single most common theme in the comments was melody. While it is hardly news that people like a good melody, this result draws our attention to the huge importance of this factor; the question of what makes a melody good is one that music theory and music psychology are still a long way from answering. Our corpus might provide a useful starting point for an exploration of this topic. Other frequent themes in respondents' comments—such as harmony, instrumentation, rhythm, and texture—also point to factors that greatly influence listeners' preferences; how they do so is, at present, largely mysterious.

4.1 Future Directions

In terms of future directions, the first avenue of exploration could be expanding the existing dataset. The survey could be re-run online, and globally, with many more participants, increasing sample size and statistical power, as well as diversifying the participant set. Instead of three songs and excerpts, many more songs and excerpts could be requested from each participant, allowing for better trend analysis *within* participants, to potentially identify different listener-types. The usefulness of this kind of data for the music-recommendation space, and associated industry applications, is clear [29].

As the corpus grows in size, the potential for using machine learning and related methods (which tend to excel with larger datasets) to analyze the data becomes more viable. An acoustical signal-analysis-based approach, using the many tools available in the field of music information retrieval, for instance, could be applied to the corpus, to determine which audio features (e.g. spectral flux, dissonance, loudness, etc.) are determinative of the favorite excerpts as compared to random controls from the same pieces. This acoustical approach could be effectively combined with a symbolic, music-theoretic approach.

In fact, even without venturing outside of the symbolic space, there is immense potential for further coding and analysis of corpus features such as scale-degree distributions, metric position, harmonic root patterns, and so forth, akin to the statistical work applied to the Rolling Stone Corpus [30, 31]. This computational approach to the corpus could be supplemented by a more humanistic, analytical approach in which more speculative and traditional analysis is conducted to attempt to understand why these particular excerpts are so powerful. For instance, given the overwhelming emphasis on pitch (melody and harmony)

in participants' comments, it would be interesting to determine the melodic, harmonic, rhythmic, and contrapuntal structures characteristic of the excerpts in the corpus; and what distinguishes a favorite excerpt from a non-favorite excerpt in the same piece.

Another possible direction for future research could be to measure the energy and complexity trajectories of the pieces in our corpus. While energy can be measured using low-level spectral features such as root-mean-square (RMS) acoustic energy, some efforts have been made to create more sophisticated predictors of perceived musical energy using combinations of features [32, 33]. Such algorithms could be applied to our corpus. Meanwhile, measuring complexity (especially in an automatic way) presents more of a challenge. Complexity—in its information-theoretic sense—is inherently subjective, since it depends on the listener's expectations, which in turn can vary widely depending on their musical experiences. Furthermore, complexity presumably depends heavily on patterns of pitch and rhythm, which cannot yet be reliably extracted from polyphonic audio [34]. For classical pieces, MIDI encodings could be used, but for popular songs, transcriptions would need to be created. Once these problems were solved, it might be possible to create measures of complexity using probabilistic models (such as Markov models); indeed, there have been interesting efforts in this direction, though they relate only to melody [35, 36].

Another intriguing area is the correlation of personality, personal values, and socio-economic data with music taste [1-6]. An expanded iteration of the survey could perhaps include a personality inventory and collect socio-economic data, building a more holistic and accurate model of music taste.

We hope that the current study has taken a small step toward advancing our understanding of peak musical experiences, and that our publicly available corpus will be useful to other researchers in this area, as we continue to answer the question: why do we like the music that we like?

5. REFERENCES

- [1] D. J. Hargreaves, C. Comber, and A. Colley, "Effects of age, gender, and training on musical preferences of British secondary school students," *Journal of Research in Music Education*, vol. 43, no. 3, pp. 242-250, 1995.
- [2] P. J. Rentfrow and S. D. Gosling, "The do re mi's of everyday life: The structure and personality correlates of music preferences," *Journal of Personality and Social Psychology*, vol. 84, no. 6, pp. 1236-1256, 2003.
- [3] S. Manolios, A. Hanjalic, and C. C. S. Liem, "The influence of personal values on music taste: Towards value-based music recommendations," *Proceedings of the 13th ACM Conference on Recommender Systems*, pp. 501-505, 2019.
- [4] T. Schäfer and C. Mehlhorn, "Can personality traits predict musical style preferences? A meta-analysis,"

- Personality and Individual Differences*, vol. 116, pp. 265-273, 2017.
- [5] A. Mohan and E. Thomas, "Effect of background music and the cultural preference to music on adolescents' task performance," *International Journal of Adolescence and Youth*, vol. 25, no. 1, pp. 562-573, 2020.
- [6] B. I. L. M. Mendis *et al.*, "Exploration of music preferences among the socioeconomic stereotypes: A cross-sectional study," *Journal of Advanced Research in Social Sciences*, vol. 4, no. 4, pp. 1-18, 2021.
- [7] A. H. Maslow, *Religions, Values, and Peak-experiences*. Columbus, OH: Ohio State University Press, 1964.
- [8] J. Whaley, J. Sloboda, and A. Gabrielsson, "Peak experiences in music," in *The Oxford Handbook of Music Psychology*, S. Hallam, I. Cross, and M. Thaut, Eds. Oxford, UK: Oxford University Press, 2009, pp. 452-61.
- [9] J. Sloboda, "Music structure and emotional response: Some empirical findings," *Psychology of Music*, vol. 19, no. 2, pp. 110-120, 1991.
- [10] A. Goldstein, "Thrills in response to music and other stimuli," *Physiological Psychology*, vol. 8, pp. 126-129, 1980.
- [11] A. J. Blood and R. J. Zatorre, "Intensely pleasurable responses to music correlate with activity in brain regions implicated in reward and emotion," *Proceedings of the National Academy of Sciences*, vol. 98, no. 20, pp. 11818-11823, 2001.
- [12] V. N. Salimpoor, M. Benovoy, K. Larcher, A. Dagher, and R. J. Zatorre, "Anatomically distinct dopamine release during anticipation and experience of peak emotion to music," *Nature Neuroscience*, vol. 14, no. 2, pp. 257-262, 2011.
- [13] O. Grewe, F. Nagel, R. Kopiez, and E. Altenmüller, "Listening to music as a re-creative process: Physiological, psychological, and psychoacoustical correlates of chills and strong emotions," *Music Perception*, vol. 24, no. 3, pp. 297-314, 2007.
- [14] S. Bannister, "A survey into the experience of musically induced chills," *Psychology of Music*, vol. 48, no. 2, pp. 297-314, 2020.
- [15] J. Panksepp, "The emotional sources of "chills" induced by music," *Music Perception*, vol. 13, no. 2, pp. 171-207, 1995.
- [16] A. Gabrielsson and S. L. Wik, "Strong experiences related to music: A descriptive system," *Musicae Scientiae*, vol. 7, no. 2, pp. 157-217, 2003.
- [17] A. Gabrielsson, "Strong experiences with music," in *Handbook of Music and Emotion*, P. Juslin and J. Sloboda, Eds. Oxford, UK: Oxford University Press, 2010, pp. 547-574.
- [18] E. Lustig, "The effect of perceived complexity and formal location on musical preference," Ph.D. dissertation, Dept. Music Theory, University of Rochester, Rochester, NY, 2021.
- [19] J. R. Landis and G. G. Koch, "The measurement of observer agreement for categorical data," *Biometrics*, vol. 33, no. 1, pp. 159-174, 1977.
- [20] P. Janata, S. T. Tomic, and J. M. Haberman, "Sensorimotor coupling in music and the psychology of the groove," *Journal of Experimental Psychology: General*, vol. 141, no. 1, pp. 54-75, 2012.
- [21] M. Witek, E. Clarke, M. Wallentin, M. Kringelbach, and P. Vuust, "Syncopation, body-movement and pleasure in groove music," *PLoS ONE*, vol. 9, no. 4, 2014.
- [22] J. A. Russell and L. F. Barrett, "Core affect, prototypical emotional episodes, and other things called emotion: Dissecting the elephant," *Journal of Personality and Social Psychology*, vol. 76, no. 5, pp. 805-819, 1999.
- [23] A. Gabrielsson and E. Lindström, "The role of structure in the musical expression of emotions," in *Handbook of Music and Emotion*, P. Juslin and J. Sloboda, Eds. Oxford, UK: Oxford University Press, 2010, pp. 367-400.
- [24] T. de Clercq, "Sections and successions in successful songs: A prototype approach to form in rock music," Ph.D. dissertation, Dept. Music Theory, University of Rochester, Rochester, NY, 2012.
- [25] P. Evans and E. Schubert, "Relationships between expressed and felt emotions in music," *Musicae Scientiae*, vol. 12, no. 1, pp. 75-99, 2008.
- [26] B. K. Hurley, P. A. Martens, and P. Janata, "Spontaneous sensorimotor coupling with multipart music," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 40, no. 4, pp. 1679-1696, 2014.
- [27] L. Meyer, *Emotion and Meaning in Music*. Chicago, IL: University of Chicago Press, 1956.
- [28] D. Huron, *Sweet Anticipation: Music and the Psychology of Expectation*. Cambridge, MA: MIT Press, 2006.
- [29] A. Agostinelli *et al.*, "MusicLM: Generating music from text," *arXiv preprint arXiv:2301.11325*, 2023.
- [30] T. de Clercq and D. Temperley, "A corpus analysis of rock harmony," *Popular Music*, vol. 30, pp. 47-70, 2011.
- [31] I. Tan, E. Lustig, and D. Temperley, "Anticipatory syncopation in rock: A corpus study," *Music Perception*, vol. 36, no. 4, pp. 353-370, 2019.
- [32] A. Zils and F. Pachet, "Extracting automatically the perceived intensity of music titles," in *Proceedings of the 6th COST-G6 Conference on Digital Audio Effects (DAFX03)*, 2003.

- [33] P. Wood and S. Semwal, "An algorithmic approach to music retrieval by emotion based on feature data," in *Proceedings of 2016 Future Technologies Conference (FTC)*, 2016, pp. 140-144.
- [34] E. Benetos, S. Dixon, Z. Duan, and S. Ewert, "Automatic music transcription: An overview," *IEEE Signal Processing Magazine*, vol. 36, no. 1, pp. 20-30, 2018.
- [35] T. Eerola, "Expectancy-violation and information-theoretic models of melodic complexity," *Empirical Musicology Review*, vol. 11, no. 1, 2016.
- [36] B. P. Gold, M. T. Pearce, E. Mas-Herrero, A. Dagher, and R. J. Zatorre, "Predictability and uncertainty in the pleasure of music: A reward for learning?," *Journal of Neuroscience*, vol. 39, no. 47, pp. 9397-9409, 2019.