# EXPLORING TONAL-DRAMATIC RELATIONSHIPS IN RICHARD WAGNER'S RING CYCLE

**Frank Zalkow, Christof Weiß, Meinard Müller**
International Audio Laboratories Erlangen, Germany
{frank.zalkow,christof.weiss,meinard.mueller}@audiolabs-erlangen.de

## ABSTRACT

Richard Wagner's cycle *Der Ring des Nibelungen*, consisting of four music dramas, constitutes a comprehensive work of high importance for Western music history. In this paper, we indicate how MIR methods can be applied to explore this large-scale work with respect to tonal properties. Our investigations are based on a data set that contains 16 audio recordings of the entire *Ring* as well as extensive annotations including measure positions, singer activities, and leitmotif regions. As a basis for the tonal analysis, we make use of common audio features, which capture local chord and scale information. Employing a cross-version approach, we show that global histogram representations can reflect certain tonal relationships in a robust way. Based on our annotations, a musicologist may easily select and compare passages associated with dramatic aspects, for example, the appearance of specific characters or the presence of particular leitmotifs. Highlighting and investigating such passages may provide insights into the role of tonality for the dramatic conception of Wagner's *Ring*. By giving various concrete examples, we indicate how our approach may open up new ways for exploring large musical corpora in an intuitive and interactive way.

## 1. INTRODUCTION

Originating in late 16th-century Florence, opera evolved as a central art form of Western culture [1]. Intended as a return to ancient Greek dramatic style, the idea of accompanied singing (*monody*) laid the ground work for two central singing styles of traditional opera: speech-like *recitatives*, which serve as a means for developing the plot, and *arias*, which emphasize the characters' feelings through cantabile melodic lines. For centuries, the structure of opera was determined by alternating such individual pieces of music, which is also known as *number opera*. In the mid-19th century, Richard Wagner developed a novel approach to operatic composition. According to his theoretical writings such as *Oper und Drama* [15], the "drama of the future" should integrate all forms of art ("Gesamtkunstwerk"). He broke with the conventions of number opera in favor of a unity of prose and music with a steady musical flow, which is often referred to as through-composed style or "endless melody" since it lacks both interruptions and exact repetitions. A central aspect of Wagner's operatic style is the frequent use of leitmotifs—short musical ideas associated with a character, a place, an item, or with emotional categories, among others.

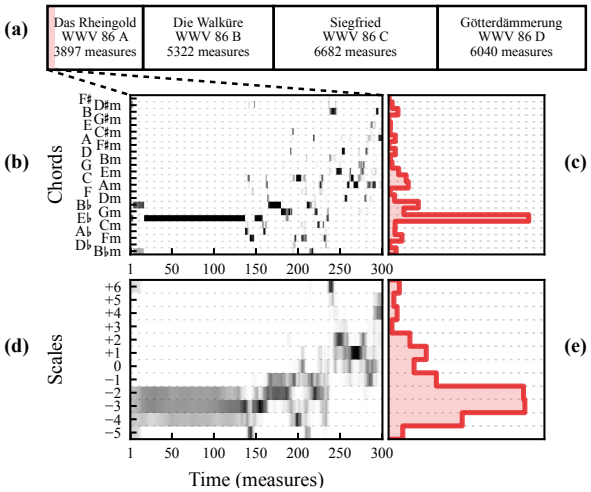


**Figure 1**. Overview of Wagner's *Ring* and schematic description of the histogram extraction. **(a)** The four parts of the *Ring* with catalogue number and length in measures. The measures under consideration are highlighted in color. **(b)** A local chord representation with time given in measures. **(c)** A histogram summarizing the chord representation. **(d)** A local scale representation. **(e)** A histogram summarizing the scale representation.

One of the most impressive realizations of these ideas is the tetralogy *Der Ring des Nibelungen*, an extensive work cycle of four music dramas created between 1848 and 1874. In Figure 1a, we show an overview of the *Ring*'s parts. A typical performance lasts 14–15 hours in total, which is demanding for listeners as well as performers. Furthermore, the through-composed form may appear less structured to the naive listener than a traditional number opera. Therefore, navigation and visualization tools are particularly useful for exploring this large-scale work. In this paper, we present such visualizations and demonstrate their benefit for musicological research.

The *Ring* has already obtained some attention in the field of Music Information Retrieval (MIR). Page et al. [9] present a toolkit for annotating musical performances in a

case study based on the *Ring*. In other studies [17, 21], Wagner's tetralogy serves as a basis for investigating the reliability of measure annotations. Concerning leitmotifs, Müllensiefen et al. [7] consider the human memory recall task. They found that the distance of chroma features relates to the perceived novelty of a leitmotif.

In this paper, we approach the *Ring* from the perspective of tonal analysis. To this end, we perform experiments on the basis of common tonal features extracted from different recordings. These audio features capture the local presence of chord [12] and scale structures [18]. Figure 1 illustrates such tonal representations with respect to major and minor triads (Figure 1b) and to diatonic scales [1] (Figure 1d) for the beginning of the first part, *Das Rheingold*. Dark values indicate a higher salience of the respective tonal structures. The prelude of this piece strongly relies on a single E♭ major triad, which corresponds to homogeneous regions in both feature sequences. In general, Wagner's music is known to present a rich vocabulary of different chord and scale types. However, though being an oversimplification, in this case study we consider only the major and minor chords as well as diatonic scales. These simple tonal structures still explain a relevant span of harmonically stable passages [19]. As a central technique in this paper, we summarize such chord and scale feature sequences in histogram representations as shown in Figures 1c/e. Visualizing these histograms may illustrate global trends in the tonal conception of the music. For computing the histograms, we select and compare passages associated with dramatic aspects, for example, the activity of certain characters or the presence of certain leitmotifs. This way, we show the ability of our visualizations to explore interesting tendencies and to study relationships between tonal and dramatic aspects within the *Ring* cycle.

In the field of MIR, histogram-based features have been extensively used for tasks such as genre recognition [13], tuning estimation [5], and other classification tasks [10, 14]. Moreover, the visualization of music pieces is an important topic in MIR. Wu and Bello [20] present an approach for visualizing musical structure. Sapp's scape plot representations [11] have come out useful for illustrating harmony analysis results in a hierarchical way. In [3], Gómez and Bonada demonstrate several visualization techniques concerning tonal aspects of musical pieces.

Typically, tonal analysis is performed on the basis of musical scores. In Section 2, we discuss why an analysis based on audio recordings may be beneficial for the *Ring* scenario. Having several recorded performances (versions) of the *Ring* allows us to employ a cross-version approach in order to stabilize the audio-based tonal representations. Konz et al. [6] show that visualizing the cross-version consistency of analysis results suppresses aspects of particular performances and, thus, emphasizes aspects relevant to the musical work in general.

Based on previously mentioned works, we present the histogram visualizations as a novel way to explore the *Ring*. The main contribution is the application of MIR techniques in an exploratory manner for highlighting interesting trends and relations within this large-scale work cycle. The remainder of the paper is structured as follows. In Section 2 we explain the data set and describe the characteristics of our annotations. Then, we shortly recapitulate the extraction of local chord and scale information and explain the histogram computation (Section 3). In the central Section 4, we discuss these histograms in detail by means of several concrete examples relating to different dramatic aspects of the *Ring*. Section 5 concludes our paper.

[1] We refer to the diatonic scales according to the respective accidentals. For example, +1 corresponds to a scale with 1♯ (G major or E minor scale) whereas −2 indicates a scale with 2♭ (B♭ major or G minor scale).

| No. | Conductor | Recording | hh:mm:ss |
|---|---|---|---|
| 1 | Barenboim | 1991–92 | 14:54:55 |
| 2 | Boulez | 1980–81 | 13:44:38 |
| 3 | Böhm | 1967–71 | 13:39:28 |
| 4 | Furtwängler | 1953 | 15:04:22 |
| 5 | Haitink | 1988–91 | 14:27:10 |
| 6 | Janowski | 1980–83 | 14:08:34 |
| 7 | Karajan | 1967–70 | 14:58:08 |
| 8 | Keilberth/Furtwängler | 1952–54 | 14:19:56 |
| 9 | Krauss | 1953 | 14:12:27 |
| 10 | Levine | 1987–89 | 15:21:52 |
| 11 | Neuhold | 1993–95 | 14:04:35 |
| 12 | Sawallisch | 1989 | 14:06:50 |
| 13 | Solti | 1958–65 | 14:36:58 |
| 14 | Swarowsky | 1968 | 14:56:34 |
| 15 | Thielemann | 2011 | 14:31:13 |
| 16 | Weigle | 2010–12 | 14:48:46 |

**Table 1**. Performances of the *Ring* used for this paper. In No. 8, Furtwängler only conducts *Die Walküre* (different from No. 4), the other parts are conducted by Keilberth.

## 2. DATA SET AND ANNOTATIONS

For an automated tonal analysis, we have to rely on a specific representation of a piece of music. Typically, musicologists perform such analyses in a manual fashion on the basis of musical scores. To automate this process, musical scores need to be accessible in a machine-readable form (symbolic data) in high quality, which is a rare case for Western classical music with a large instrumentation. To overcome this problem, Optical Music Recognition (OMR) techniques are usually employed, which cannot provide satisfactory results in many situations [2] so that time-consuming manual correction steps are required. As an alternative, tonal analysis can be performed on the basis of audio recordings, at least to a certain extent [12, 18]. For the *Ring*, for example, a high number of CD recordings are easily available. In our experiments, we use 16 different performances comprising nearly 232 hours of audio, see Table 1. To compare and combine analysis results obtained from different representations, a link between the representations' time axes is beneficial. For instance, we can use the positions of measure boundaries, as specified by the score, in the recordings. To this end, several students with a strong practical experience in Western classical music manually annotated these positions for three performances [17]. By means of synchronization techniques [8, Chap. 3], we jointly transferred the manual annotations from the three performances to all other performances. See [21] for details and an evaluation regarding

the joint transfer. The measure positions constitute suitable reference points for a cross-version analysis on the measure level since different performances can be related to each other using a musical time axis. In total, the *Ring* encompasses 21941 measures, including pickup measures.

Moreover, the measure positions enable the transfer of semantic annotations from a musical time axis to the individual performances and vice versa. In our scenario, we are interested in dramatic aspects of the *Ring*'s plot. For example, we annotated the regions where different characters are singing. In particular, we specified the start and end of the singing voice regions as well as the corresponding character according to the verses of the libretto. If a verse is interrupted by a whole measure or more, the annotation is split accordingly. In total, our annotations comprise 6792 singing voice regions.

Beyond this, the presence of certain leitmotifs is of high interest for studying the dramatic conception. Motivated by this, a musicologist annotated the occurrences of the leitmotifs by listening to a recording and analyzing the sheet music. The definition of leitmotifs used for this work relies on a guide by Julius Burghold from 1913 [16], which serves as a reference both for the names and the musical shape of the motifs. The annotations comprise start and end positions as well as the corresponding name for each occurrence of a leitmotif. For example, within *Siegfried* 2632 leitmotif occurrence regions have been identified.

## 3. HISTOGRAM COMPUTATION

In this section, we summarize the computation of local chord and scale representations, describe our histogram approach, and show why a cross-version strategy is beneficial. To capture tonal characteristics in audio recordings, we first compute normalized chroma features [8, Chap. 3], which represent the energy within the twelve chromatic pitch classes over time. Employing a common template matching strategy, we locally compare the chroma feature sequence with binary templates corresponding to chords [12] or scales [18]. By means of normalizing, we can interpret the results as probability for the occurrences of the particular chords or scales. Since these tonal structures relate to a time span of several seconds, we suitably smooth the chromagram before applying template matching. To obtain musically meaningful windows, we make use of the measure annotations to obtain performance-independent window sizes $w \in \mathbb{N}$ specified in measures, rather than seconds. Our experiments showed that $w = 4$ for the chord analysis and $w = 12$ for the scale analysis provides meaningful visualizations. From a traditional music theory perspective, a window size of 4 measures for analyzing chords does not make sense. Indeed, chords lasting for such a long time rarely occur in Wagner's music. However, such a parameter setting leads to visualizations, which appear more structured, e.g. emphasizing tonic chords. One reason is that many chords in a progression share common notes and, thus, often stabilize the result for the respective tonic chord. Since we use a centered window view and a hop size of one measure, we obtain one feature vector for each measure. For emphasizing the locally salient structures, we apply an exponential post-processing step similar to the softmax function.

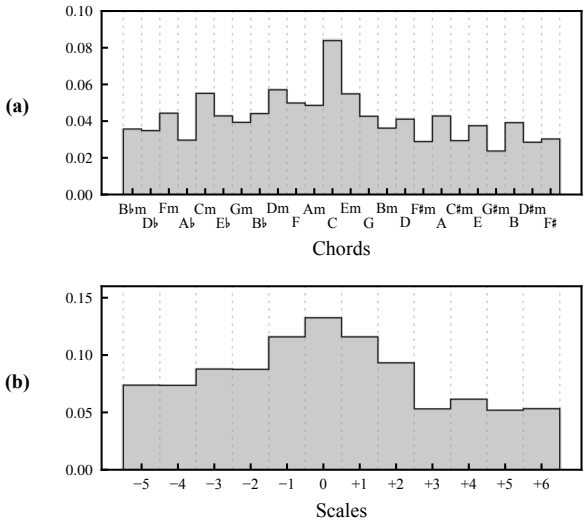


**Figure 2**. Analysis for the complete *Ring* cycle. **(a)** Chord histogram. **(b)** Scale histogram.

With this strategy, we compute an analysis matrix $\mathcal{A} \in \mathbb{R}^{D \times M}$ for a given performance, where $M \in \mathbb{N}$ denotes the number of measures and $D \in \mathbb{N}$ the feature dimension (with $D = 24$ for chords and $D = 12$ for scales). Figures 1b/d show such matrices. Based on this matrix $\mathcal{A}$, we calculate a histogram $\mathbf{h} \in \mathbb{R}^D$ by averaging over all measures:

$$\mathbf{h}(d) = \frac{1}{M} \sum_{m=1}^{M} \mathcal{A}(d, m) \tag{1}$$

with $d \in [1 : D] := \{1, 2, \ldots, D\}$. Figure 2 shows histograms for the complete *Ring*. The bar heights correspond to the presence of the chords or scales averaged over all measures. The distributions are rather flat, which indicates that Wagner seems to use the full range of chords and keys for tonally shaping his tetralogy. Nevertheless, we observe a stronger presence of the C major chord as well as the scales 0, +1, and −1 indicating that tonal regions with few accidentals seem to be slightly more prominent. Furthermore, we find a small trend towards flat key regions (left half) compared to sharp key regions (right half).

At this point, we may wonder about the reliability of these results and the histograms' dependency on a specific version. For example, Figure 3a shows two chord histograms, which are computed for two different performances. Even though the global trends seem to be consistent among the two versions, one can observe some deviations. For example, the E minor triad seems to be one of the most prominent chords in the Boulez version (blue), whereas this chord is less important in the Furtwängler version (red). Such performance-specific characteristics may come into play for a variety of reasons. For example, performances may exhibit a different dynamic balance of the instruments and singers. Furthermore, different recording conditions may suppress or enhance certain frequencies.

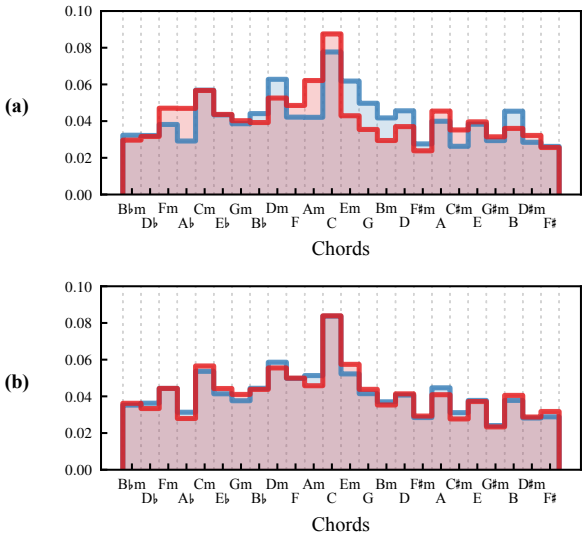


**Figure 3**. Comparison of single-version histograms and cross-version histograms for the complete *Ring*. **(a)** Chord histograms for the performances conducted by Boulez (No. 2) in blue and Furtwängler (No. 4) in red. **(b)** Chord histograms with cross-version approach. The blue and the red histogram each correspond to eight different versions.

To attenuate the version-specific aspects, we introduce a cross-version approach similar to [6]. Having $P \in \mathbb{N}$ performances, we compute an analysis matrix $\mathcal{A}_p$ for each of these versions $p \in [1:P]$. From these matrices, we derive a cross-version analysis matrix simply by averaging over all performances

$$\mathcal{A}^{\mathrm{cv}}(d,m) = \frac{1}{P}\sum_{p=1}^{P}\mathcal{A}_p(d,m) \qquad (2)$$

with $d \in [1:D]$ and $m \in [1:M]$. Finally, we average over all measures of $\mathcal{A}^{\mathrm{cv}}$ to obtain a histogram as in Eq. 1.

Figure 3b shows two histograms computed with our cross-version approach, each for $P = 8$ different performances. [2] Note that the two cross-version histograms are more similar to each other than the two single-version histograms in Figure 3a. This indicates that the cross-version approach stabilizes the results. Furthermore, characteristic trends and peaks of the histograms are retained, which shows that the averaging procedure does not smooth out interesting details. In general, the cross-version approach enhances work-related aspects and suppresses performance-specific artifacts. In the following, all histograms are computed in a cross-version fashion using $P = 16$ performances. Also, the histograms in Figure 2 were computed in this way.

## 4. EXPLORATION

We now want to show the potential of the introduced histograms for exploring relationships between dramatic aspects of the *Ring* and its tonal organization. To this end, we

[2] Performances Nos. 1, 2, 5, 9, 11–13, 15 correspond to blue and performances Nos. 3, 4, 6–8, 10, 14, 16 correspond to red.

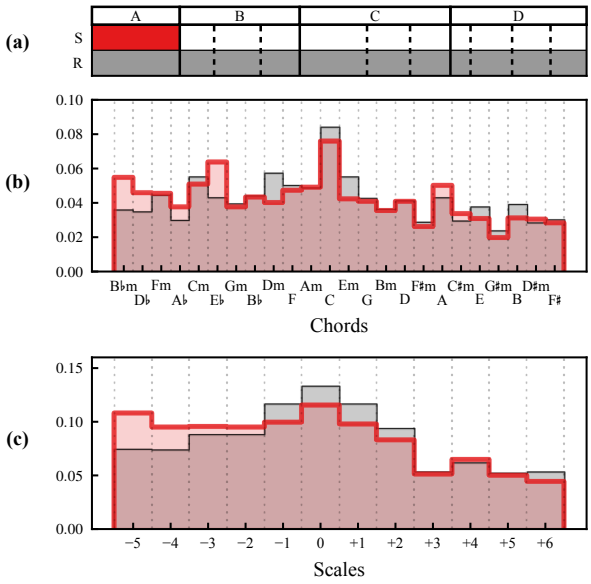


**Figure 4**. Comparison of *Das Rheingold* to the complete *Ring*. **(a)** Activations for the *Ring* (reference, 21941 measures, gray) and *Das Rheingold* (selection, 3897 measures, red). Dotted lines correspond to the beginning of a new act. **(b)** Chord histogram. **(c)** Scale histogram.

apply a selection procedure with respect to different criteria. As examples of such criteria, we consider the role of a single part with respect to the full cycle, the behavior of singing and instrumental regions, as well as the activity of specific characters. Furthermore, we take the occurrence of certain leitmotifs into account. For all examples, we select the measures or passages fulfilling these criteria while discarding the others before computing the histograms.

As a first scenario, we looked at the role of individual parts of the *Ring*. For example, we compared the first part, *Das Rheingold*, to the complete cycle. In Figure 4a, we show a compact overview that illustrates the investigated passages with the complete cycle as reference (R) in gray and the first part as selection (S) in red. In the following, we refer to such illustrations as activation diagrams. Using this color scheme, Figure 4b shows two chord histograms with the reference histogram in gray and the selection histogram in red. In these histograms, we observe rather flat distributions. Regarding the individual chords, we observe a high presence of the C major chord in *Das Rheingold* (selection). This coincides with the general distribution in the full *Ring* (reference). In contrast, the enhanced presence of the E♭ major chord in the selection deviates from the shape of the reference histogram. An important reason for this peak may be the prelude comprising 136 measures with a constant harmony, an E♭ major triad. Figure 4c shows the corresponding histograms for the scales. The comparison of these histograms indicates a trend towards scales with flats (♭) in their key signatures (–2 to –5) in the selection. These observations suggest that the individual parts of the *Ring* may indeed exhibit a characteristic tonal shape.

As a second example, we examine the characteristics of passages involving instrumental passages and singing

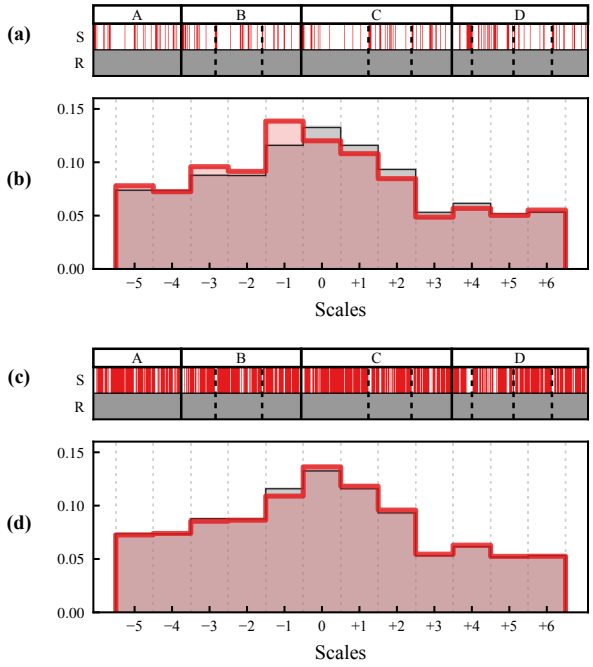


**Figure 5**. Comparison of instrumental passages to all others. **(a)** Activations for the *Ring* (reference, 21941 measures, gray) and the instrumental passages (selection, 5103 measures, red). **(b)** Scale histogram. **(c)** Activations for the *Ring* (reference, 21941 measures, gray) and passages involving singing (selection, 16838 measures, red). **(d)** Scale histogram.

activity. In contrast to the previous experiment, our selections now comprise several local passages instead of one contiguous region. To obtain two disjoint selections, we assign each measure to either the instrumental or the singing selection where measures with partial singing were assigned to the singing selection. Figures 5a/c display the corresponding activation diagrams. The aggregated lengths of both subsets are in a proportion of roughly 1 : 3 (5103 vs. 16838 measures). In the scale histograms in Figures 5b/d, the selections show a high similarity to the full cycle's histogram even though the respective passages do not overlap. This indicates that, globally speaking, there is no substantial tonal difference between these selections. One may wonder if this similarity is a trivial observation. As we will show in our next example, we observe different shapes when we investigate the singing activity of specific characters.

As the third scenario, we examine the behavior for two groups of characters, *gods* and *mortals*, which constitute central categories in the *Ring*'s plot. [3] The activation diagram in Figure 6a shows the singing activities of gods (in blue) and mortals (in red). Instead using a reference, like the complete *Ring* as before, we now compare two selections based on different selection criteria. In *Das Rheingold*, only gods are active whereas in the final *Götterdämmerung*, gods are barely active (only the Norns in the prologue). Thus, we observe some kind of global trend from gods towards mortals over the course of the *Ring*. Figures 6b/c show the corresponding chord and scale histograms. The mortals' scale histogram exhibits a slight trend towards scales with few accidentals. In contrast, the gods' histogram shows a somewhat higher presence of outlying scales. For musicologists, such observations could be a starting point for relating tonal characteristics to the interpretation of the drama. As an exemplary hypothesis, the prominence of far-off scales might be associated with the gods living far-off our human world.

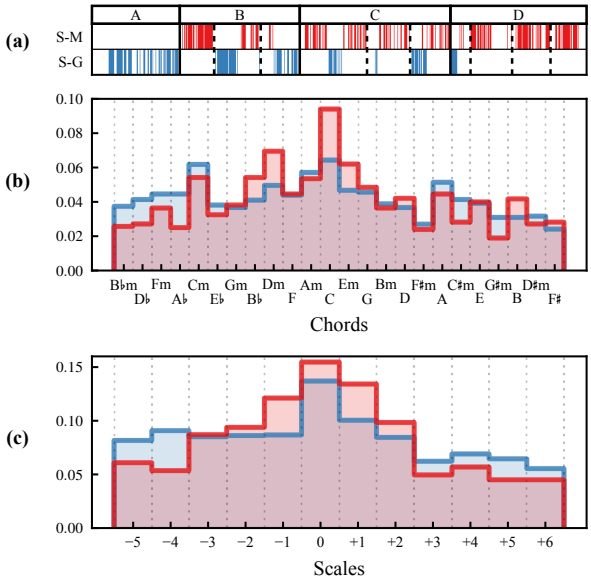


**Figure 6**. Comparison of gods and mortals singing. **(a)** Activations for passages with mortals singing (S-M, 6409 measures, red) and passages with gods singing (S-G, 4203 measures, blue). **(b)** Chord histogram. **(c)** Scale histogram.

As a final selection criterion, we focus on the occurrence of leitmotifs, which constitute a central dramatic element in the *Ring*. Even though Wagner did not invent this technique and never used the term "leitmotif" personally, the extensive usage of such motifs makes the *Ring* a prominent example of a realization of this concept. We now indicate how one may explore the tonal characteristics during the occurrences of certain leitmotifs. As a first example, we consider the "Valhalla motif," which refers to the castle of the gods and is frequently used over the course of the tetralogy. Figures 7a/b show an activation diagram and a chord histogram with *Das Rheingold* as reference and all regions in this part with the Valhalla motif as selection. For this motif, we notice a high peak for the D♭ major chord, in contrast to the flat shape of the reference histogram. When we examine this motif over the course of *Die Walküre*, we observe a different trend (see Figures 7c/d). For this part, the histogram exhibits a high peak for the E major chord whereas the D♭ major chord has only a slight peak. These

[3] For this exemplary scenario, we classified as gods: Wotan, Fricka, Freia, Donner, Froh, Erda, Loge, and the Norns. Characters classified as mortals are Siegmund, Sieglinde, Siegfried, Hunding, Gunther, Gutrune, Hagen, as well as the male and female choirs. Even though there are border cases such as the demigod Loge or Siegmund and Sieglinde, Wotans children, we consider this a meaningful categorization. We do not take into account other categories such as the Valkyries, or the Nibelungs.

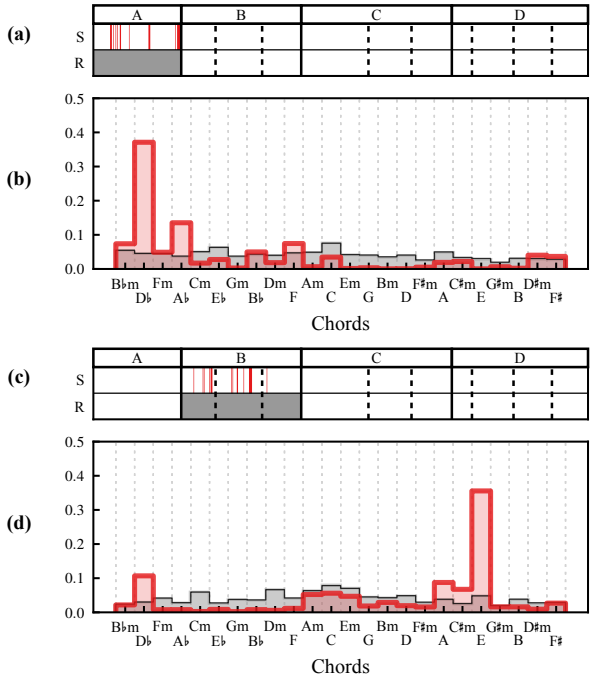


**Figure 7**. Passages containing the Valhalla motif. **(a)** Activations for *Das Rheingold* (reference, 3897 measures, gray) and passages containing the Valhalla motif therein (selection, 176 measures, red). **(b)** Chord histogram. **(c)** Activations for *Die Walküre* (reference, 5322 measures, gray) and passages containing the Valhalla motif therein (selection, 124 measures, red). **(d)** Chord histogram.

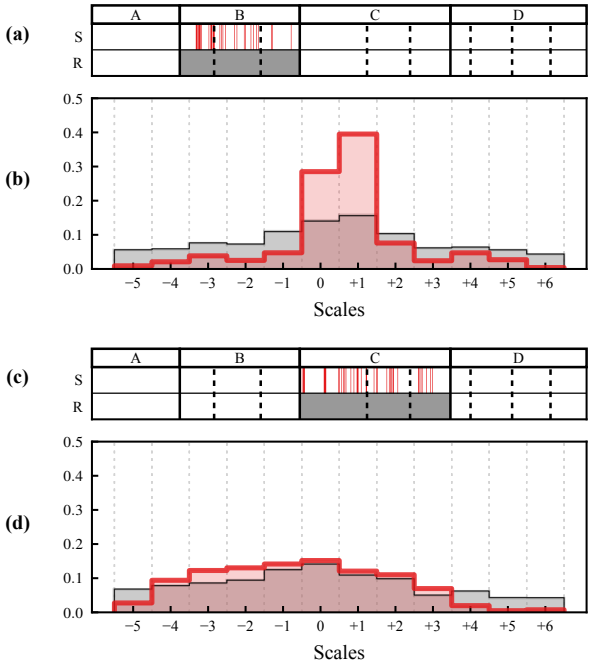


**Figure 8**. Passages containing the sword motif. **(a)** Activations for *Die Walküre* (reference, 5322 measures, gray) and passages containing the sword motif therein (selection, 200 measures, red). **(b)** Scale histogram. **(c)** Activations for *Siegfried* (reference, 6682 measures, gray) and passages containing the sword motif therein (selection, 203 measures, red). **(d)** Scale histogram.

observations indicate that the Valhalla motif tends to appear in a specific tonal context. However, the corresponding chord strongly differs between the parts of the *Ring*. The correlation of the Valhalla motif to the D♭ major and E major chords, respectively, is a known fact in musicology [4, p. 172]. It is promising that automated methods can confirm such observations.

As a second leitmotif example, we investigate the occurrences of the "sword motif." Figure 8 shows activation diagrams and scale histograms for this motif within *Die Walküre* and *Siegfried*. In *Die Walküre*, the motif has a clear tendency towards the region of 0 (C major/A minor) and +1 (G major/E minor). The situation is completely different in *Siegfried*, where the scale distribution is rather flat, with a slight trend towards flat scale regions. One reason might be the integration of this motif into a more complex tonal conception in *Siegfried* compared to *Die Walküre*.

## 5. CONCLUSIONS AND OUTLOOK

In this paper we demonstrated how existing MIR techniques such as tonal audio features and global histograms can be applied in a complex music scenario. Regarding audio-based analysis, we showed that a cross-version approach is able to enhance work-related properties while suppressing performance-specific details and, therefore, stabilizes the analysis results compared to a single-version approach. For exploring relationships between dramatic aspects and the tonal organization of Wagner's *Ring* cycle, we presented compact visualizations of tonal features based on global histograms with several selection criteria. We showed that these visualizations can provide interesting insights into large-scale works such as Wagner's tetralogy. Investigating a small selection of examples, we found that the *Ring*'s parts may each exhibit an individual tonal shape. Furthermore, the histograms indicated that character groups can have different tonal preferences. Finally, we showed that leitmotifs can have specific tonal connotations. In general, using global histograms exhibits some limitations since we cannot address many fine-granular issues with this method. Nevertheless, we showed the benefit of such visualizations for highlighting interesting trends and relations. Beyond the analyses shown in this paper, many more filtering criteria could be of interest in this complex scenario. Concerning subsections of the *Ring*, individual acts or scenes could be analyzed. Regarding the musical parts, different character groups or even individual singers could be considered as well as the use of certain instruments or instrument families. Finally, the enormous number of leitmotifs with several thousand occurrences and their complex relationships lead to a vast amount of possible selections. Investigating the *Ring* with respect to such aspects could allow musicologists to confirm or adjust their hypotheses or be inspired to create new ones and, thus, might have potential for musicological research.

**Acknowledgments:** We thank Julia Zalkow, Vlora Arifi-Müller, and all students for their assistance in generating the annotations and preparing the data. This work has been supported by the German Research Foundation (DFG MU 2686/7-1). The International Audio Laboratories Erlangen are a joint institution of the Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU) and Fraunhofer Institut für Integrierte Schaltungen IIS.

## 6. REFERENCES

[1] Howard Mayer Brown, Ellen Rosand, Reinhard Strohm, Roger Parker, Arnold Whittall, Roger Savage, and Barry Millington. Opera. In Stanley Sadie, editor, *The New Grove Dictionary of Music and Musicians*, pages 416–471. Macmillian Publishers, London, UK, 2nd edition, 2001.

[2] Donald Byrd and Jakob G. Simonsen. Towards a standard testbed for optical music recognition: Definitions, metrics, and page images. *Journal of New Music Research*, 44(3):169–195, 2015.

[3] Emilia Gómez and Jordi Bonada. Tonality visualization of polyphonic audio. In *Proceedings of the International Computer Music Conference (ICMC)*, pages 57–60, Barcelona, Spain, 2005.

[4] Tobias Janz. *Klangdramaturgie: Studien zur theatralen Orchesterkomposition in Wagners „Ring des Nibelungen“*, volume 2 of *Wagner in der Diskussion*. Königshausen & Neumann, Würzburg, Germany 2006.

[5] Gopala Krishna Koduri, Joan Serrà, and Xavier Serra. Characterization of intonation in carnatic music by parametrizing pitch histograms. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 199–204, Porto, Portugal, 2012.

[6] Verena Konz, Meinard Müller, and Rainer Kleinertz. A cross-version chord labelling approach for exploring harmonic structures—a case study on Beethoven's Appassionata. *Journal of New Music Research*, 42(1):61–77, 2013.

[7] Daniel Müllensiefen, David Baker, Christophe Rhodes, Tim Crawford, and Laurence Dreyfus. Recognition of leitmotives in Richard Wagner's music: An item response theory approach. In *Analysis of Large and Complex Data*, pages 473–483. Springer, Cham, Switzerland, 2016.

[8] Meinard Müller. *Fundamentals of Music Processing*. Springer, Cham, Switzerland, 2015.

[9] Kevin Page, Terhi Nurmikko-Fuller, Carolin Rindfleisch, Richard Lewis, Laurence Dreyfus, and David De Roure. A Toolkit for Live Annotation of Opera Performance: Experiences Capturing Wagner's Ring Cycle. In *Proceedings of the International Conference on Music Information Retrieval (ISMIR)*, pages 211–217, Málaga, Spain, 2015.

[10] Hendrik Purwins, Benjamin Blankertz, Klaus Obermayer, and Guido Dornhege. Scale degree profiles from audio investigated with machine learning. In *Proceedings of the 116th Audio Engineering Society (AES) Convention*, Berlin, Germany, 2004.

[11] Craig Stuart Sapp. Visual hierarchical key analysis. *ACM Computers in Entertainment*, 3(4):1–19, 2005.

[12] Alexander Sheh and Daniel P.W. Ellis. Chord segmentation and recognition using EM-trained hidden Markov models. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 185–191, Baltimore, USA, 2003.

[13] George Tzanetakis and Perry Cook. Musical genre classification of audio signals. *IEEE Transactions on Speech and Audio Processing*, 10(5):293–302, 2002.

[14] George Tzanetakis, Andrey Ermolinskyi, and Perry Cook. Pitch histograms in audio and symbolic music information retrieval. *Journal of New Music Research*, 32(2):143–152, 2003.

[15] Richard Wagner. *Oper und Drama*. Klaus Kropfinger (Ed.). Reclam, Ditzingen, Germany, 1986.

[16] Richard Wagner. *Der Ring des Nibelungen. Vollständiger Text mit Notentafeln der Leitmotive*. Julius Burghold (Ed.). Schott Music, Mainz, Germany, 2013.

[17] Christof Weiß, Vlora Arifi-Müller, Thomas Prätzlich, Rainer Kleinertz, and Meinard Müller. Analyzing measure annotations for western classical music recordings. In *Proceedings of the International Conference on Music Information Retrieval (ISMIR)*, pages 517–523, New York, USA, 2016.

[18] Christof Weiß and Julian Habryka. Chroma-based scale matching for audio tonality analysis. In *Proceedings of the 9th Conference on Interdisciplinary Musicology (CIM)*, pages 168–173, Berlin, Germany, 2014.

[19] Christof Weiß, Rainer Kleinertz and Meinard Müller. Möglichkeiten der computergestützten Erkennung und Visualisierung harmonischer Strukturen – eine Fallstudie zu Richard Wagners *Die Walküre*. In *Bericht zur Jahrestagung der Gesellschaft für Musikforschung (GfM)*, Mainz, Germany, 2016.

[20] Ho-Hsiang Wu and Juan P. Bello. Audio-based music visualization for music structure analysis. In *Proceedings of the Sound and Music Computing Conference (SMC)*, pages 496–501, Barcelona, Spain, 2010.

[21] Frank Zalkow, Christof Weiß, Thomas Prätzlich, Vlora Arifi-Müller, and Meinard Müller. A multi-version approach for transferring measure annotations between music recordings. In *Proceedings of the AES International Conference on Semantic Audio*, pages 148–155, Erlangen, Germany, 2017.