CorpusVis: Visual Analysis of Digital Sheet Music Collections

Matthias Miller[1] , Julius Rauscher[1], , Daniel A. Keim[1] , and Mennatallah El-Assady[2]

[1]University of Konstanz, Germany   [2]ETH AI Center, Zürich, Switzerland
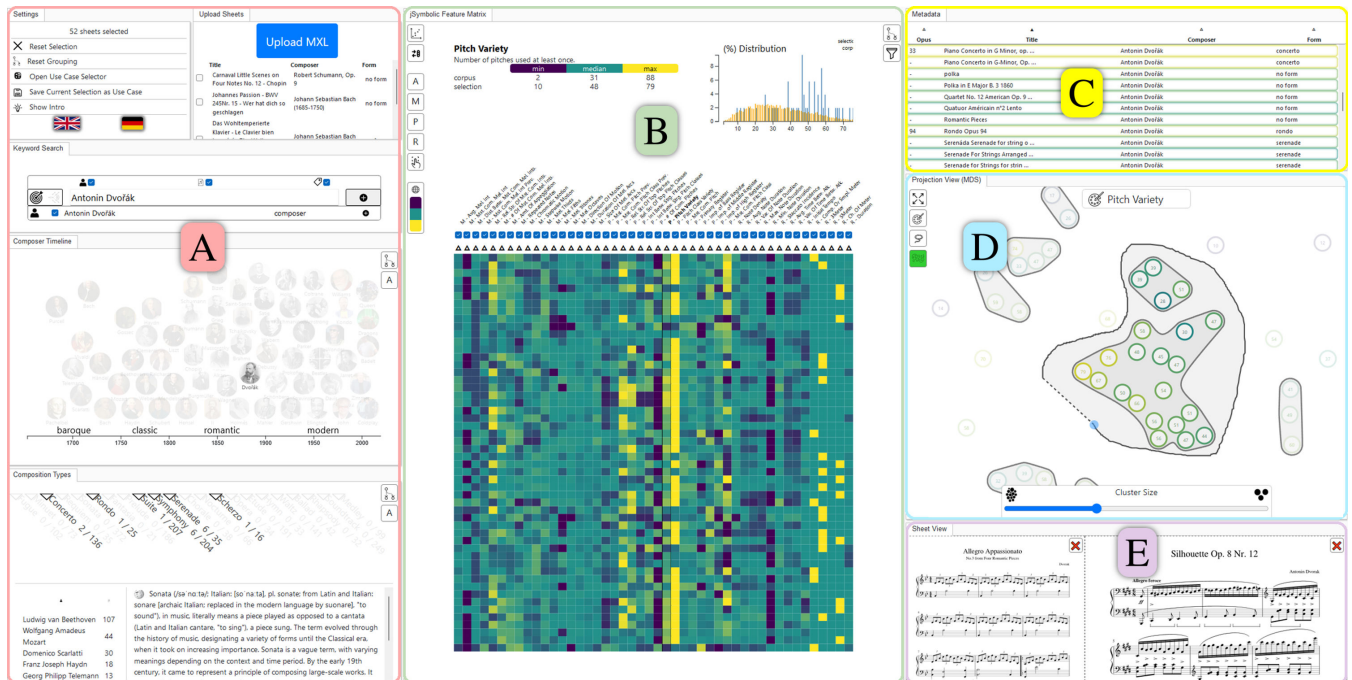
Figure 1: *Our visual interactive workspace supports the analysis of sheet music collections through visualizations that are connected through linking and brushing. After filtering a sheet music corpus* A *, music analysts can perform exploration, comparison, and detection tasks. The feature matrix* B *provides detailed information about low-level characteristics for musical compositions. The metadata table* C *enables browsing through the titles, composers, and composition forms. An MDS projection view* D *helps to identify similar compositions or even clusters based on selected features. The sheet view* E *allows analysts to view compositions using the familiar notation. Domain experts can use CorpusVis to confirm and generate hypotheses and detect interesting patterns between composers and composition types.*

Abstract

Manually investigating sheet music collections is challenging for music analysts due to the magnitude and complexity of underlying features, structures, and contextual information. However, applying sophisticated algorithmic methods would require advanced technical expertise that analysts do not necessarily have. Bridging this gap, we contribute CorpusVis, an interactive visual workspace, enabling scalable and multi-faceted analysis. Our proposed visual analytics dashboard provides access to computational methods, generating varying perspectives on the same data. The proposed application uses metadata including composers, type, epoch, and low-level features, such as pitch, melody, and rhythm. To evaluate our approach, we conducted a pair-analytics study with nine participants. The qualitative results show that CorpusVis supports users in performing exploratory and confirmatory analysis, leading them to new insights and findings. In addition, based on three exemplary workflows, we demonstrate how to apply our approach to different tasks, such as exploring musical features or comparing composers.

CCS Concepts

• Human-centered computing → Visualization techniques; Visual analytics; Visualization systems and tools; • Information systems → Content analysis and feature selection; Document structure;

1. Introduction

Investigating music collections are relevant for both domain experts such as musicologists and music theorists and usual music consumers [WJVR16]. Predominantly, the audio format has been the starting point for many applications including music classification [FLTZ11], personalized music recommendation [BHF*13], structure analysis [PMK10], and the generation of music playlists [STOH21]. In contrast, we are convinced that the investigation of sheet music collections based on digital symbolic formats such as MusicXML [Goo01], MIDI, Humdrum [Sap15], and Lilypond [NN03] have received less attention compared to audio.

While services such as sheet music recommendation offered by, e.g., OKTAV [OKT21] leverage musical features and individual preferences to provide sheet music suggestions for piano players, it remains unclear how such recommendations are computed. To understand which musical features (e.g., structure or melody) reveal differences and similarities, computational approaches for sheet music analysis have been proposed in the community, supporting various tasks such as genre and composer classification [ANC17]. Besides close reading of sheet music as a typical task in musicology, distant reading is also an essential task in the digital humanities [JFCS15]. Close reading supports musicologists to investigate musical compositions on a detailed level, abstract visualization contain the potential to provide an overview over a large set of musical compositions within their context, thus enabling distant reading. While approaches that solely rely on algorithmic or computational approaches mitigate the tedium of manual analysis of a dataset done by human analysts, they often are not accessible to a broader user group. Still, they are limited to music analysts that have programming skills. To address this gap, tailored visualizations combined with user interaction can be employed to increase access to music analysis methods for a broader user group while taking the human analyst in the analysis loop, which is a fundamental aspect for knowledge generation in visual analytics [SSS*14]. An essential aspect of the understanding process is critical thinking by the analyst about the subject, which was introduced by Bradley et al. as *slow analytics* [BEAC*18].

The availability of sheet music datasets such as KernScores [Sap05] or MuseScore [Mus11] have the potential for analysis at different scales, benefiting from using projection techniques based on underlying features [MCF18]. While MuseScore contains more than a million compositions, the quality of the uploaded material varies from content that is provided by users who ensure that provided content is faithful to the original compositions to pieces that suffer from inaccuracy or incompleteness. Therefore, it is challenging to maintain a high quality of the underlying data for the analysis when considering the full dataset, requiring expensive data cleaning steps before performing analysis tasks on it. Manual curation or selection processes are required to set the focus on a particular subset that is either provided by users, that only provides high-quality content. Alternatively, analysts can manually view each composition that shall be part of a curated dataset, which is, of course, quite a tedious process. The research field of optical music recognition provide approaches to convert printed sheet music into symbolic formats such as MusicXML. Yet, it needs further improvements until musicologists trust these automatic results [SF20]. Enabling analysts to influence the data foundation is crucial for the analysis.

Besides data quantity, increasing the quality of the data is crucial for effective analysis. This includes a targeted selection of representative samples. For example, a basic issue in humanities research is striving to avoid the exclusion of marginalized positions, a typical risk of data colonialism. For instance, famous classical composers such as Johannes Sebastian Bach, Wolfgang Amadeus Mozart, Joseph Haydn, or Ludwig van Beethoven are often a more prominent analysis subject in musicology than less known composers such as Muzio Clementi, Domenico Scarlatti, or Francois-Joseph Gossec. This marginalization of composers is reaffirmed by the vast amount and duplicates of famous compositions in datasets such as MuseScore. To enable a large-scale analysis of sheet music collections, a corpus-level overview is essential. This has the power to enable analysts to audit and refine corpora, as well as compare music sheets on numerous features. Applying abstract visualization techniques to sheet music corpora at a larger scale has the potential to provide insights about the work of composers or typical differences between compositions types without the need to manually analyze every detail of single data items.

Driven by the need for a corpus-level music analysis technique, in this work, we address the research question: *How to support music analysts to explore, investigate, and compare sheet music collections based on metadata and low-level features using interactive visualization?* Our aim is to provide them with a multiscale and multi-perspective bird's-eye view on sheet music corpora. We thus designed an interactive visual workspace that provides multiple tailored analysis components in an inter-linked dashboard. The visual analysis is supported by computational methods, such as clustering, to aid in pattern finding. Analysts can openly explore the underlying corpus, investigate a set of pre-configured use cases, or verify hypotheses through crafting their own analysis workflow.

Contributions – This work contributes a problem characterization addressing visualization requirements with regard to the analysis of sheet music corpora. We provide a list of relevant data and task characteristics, as well as a description of target audiences. A major contribution is the Visual Analysis Workspace for sheet music collections through a combination of multiple components that are seamlessly connected through linking and brushing. We conducted a qualitative evaluation to assess the applicability of our approach and provide details about its benefits and drawbacks. Finally, we discuss open research opportunities to inspire interdisciplinary collaboration at the interface of visualization and musicology.

2. Related Work

The field of musicology is a wide research area that covers heterogeneous research questions and challenges. The subfield *Visual Musicology* at the interface of musicology and visualization research as introduced by Miller et al. illustrates the vast opportunities for which visualization could be applied to support domain-related, scientific issues [MSK*19]. Their framework highlights the potential of interactive visualization to perform analysis tasks, including information retrieval, exploration, and comparison. We use the *visual musicology graph* to classify the work presented in this paper accordingly. Specifically, we focus on *Structural Features* and *Meta-Information* of sheet music and the visualization tasks *Overview/Summarization*, *Navigation/Exploration*, *Clustering*, *Comparison*, and *Details on Demand* within the domains *Theory & Analysis* and *History*.

Visual Analysis of Music Collections – The visual investigation of musical data collections has already received attention by visualization researchers. Khulusi et al. created a web-based interface called "musiXplora" as part of a digitization project that enables humanities researchers to investigate musicological data such as meta-information (e.g., gender, religious denomination, profession, institutions) of musicians and instrument makers [KKFJ20]. For instance, this interactive prototype facilitates information retrieval and browsing about the life and work of different composers. Similarly, Jänicke et al. designed an interactive visualization tool for the interactive profiling of musicians and their relationships based on their meta-information [JFS16]. They did not consider specific features about the compositional work but rather focus on the contextual information about the life, profession, and instruments they mastered. Pampalk used SOMs to visualize and classify the genre of music data as "Islands of Music" based on audio-extracted features such as the loudness information from audio files [Pam01]. Chen and Putz designed an interactive UI for browsing and organizing music collection allowing to listen and visually explore the similarity of musical pieces based on low-level features that reflect the genre and style of pieces from different composers [CB09]. Weiß et al. show how exact tonal features (pitch information) and metadata from audio data help to detect style changes over the different epochs from baroque to modern music [WMDM18]. Georges and Nguyen use a dataset of 500 classical composers to visually analyze different epochs using dendrograms and MDS projection techniques [GN19]. By that, they demonstrate how abstract visualization methods support the comparison of composers [BMHC16].

Feature Extraction from Music Data – Music is available and can be stored based on different formats such as symbolic (e.g., MusicXML [Goo01]) or audio data (e.g., mp3). Depending on this format, different features can be extracted that can afterwards be used for different tasks such as music recommendation [Sch19]. Similar to extracting features from audio signal data (e.g., spectral information) for further processing [LWSF20], it is possible to extract low-level features from symbolic music [Cc21]. For example, music21 [Cc21] is a Python library that enables programmers to extract jSymbolic features from different symbolic sheet music formats including MusicXML. McKay and Fujinaga discuss how such features can be used for MIR research tasks [MCF18] such as composer classification as done by Verma and Thickstun leveraging convolutional networks [VT19]. McKay leverages audio and symbolic music data sources to perform music classification and the creation of music information retrieval tools using his own jMIR software suite that can even consider contextual/cultural metadata [MF09]. Corrêa and Rodrigues published a survey about music genre classification that reveals that such tasks primarily use automatic algorithms without using advanced visualization techniques [CR16]. Merely computational approaches do not enable music analysts to explore low-level features, making it impossible to step into the analysis process to get a better overview over the inner workings of the applied algorithms.

Visualization of Collections in the Digital Humanities – Besides musicology, there exist other scientific issues within digital humanities that have already been addressed by information visualization researchers such as the visual analysis of poems [MLCM16]. Jänicke et al. discuss many digital humanities projects that require text pro-

cessing methods [JFCS17] while Kirschenbaum also confirms the essentiality of text for humanities research ranging from close via not-reading to distant reading [Kir07]. For example, Bludeau et al. implemented a web-based prototype to enable the visual investigation of literature and handwritten notes from *Fontane's Handbibliothek* [BBBD20]. For instance, the field of text analysis deals with comparison tasks for plagiarism detection by identifying typical features for specific authors [GMB*14]. Keim and Oelke use textual features to create a literature fingerprinting visualization that facilitates the analysis of whole books [KO07]. We argue that existing visualization methods that have been successfully applied to similar endeavors in related research areas enclose the potential for methodology transfer [MSK*19]. Consequently, we consider it to be useful to get inspired by existing visual techniques to address unsolved challenges such as the visual analysis of sheet music collections.

Research Gap – While there are several projects within this research field, they are often limited to audio data only [CB09, WMDM18]. Thus, we argue that the analysis of music based on symbolic formats has received less attention. We argue this shift towards audio data could be due to larger user groups of music applications such as (e.g., Spotify) which benefit from effective music recommendation methods. We see the reason for this unbalance in the smaller group size of sheet music consumers, which is much smaller. Liem et al. discuss the issue of focusing on audio signal only and argue for multimodal and user-centered analysis strategies to improve the accessibility of digital music data [LME*11, MGD11]. Nevertheless, the increased use of digital symbolic music formats led to services such as OKTAV [OKT21] which provides sheet music recommendations based on symbolic features. While automatic solutions have their application areas, employing visual methods allows analysts to step into the analysis process, helping them to gain new insights or generate and confirm hypotheses, which is also known as human-in-the-loop approaches [SSS*14]. To our knowledge, there is no previous work that readily supports the analysis of sheet music corpora, facilitating the visual comparison of metadata and low-level features.

3. Problem Description

Our primary objective is to provide an *interactive visual workspace* to support the investigation and exploration of sheet music collections. The target audience primarily comprises musicologists and music theorists. We assume that music librarians and Music Information Retrieval (MIR) researchers could benefit from our work as well. For this reason, we gathered information about the data and task requirements by conducting initial expert interviews and consulted existing literature.

Feature Characteristics – When analyzing musical compositions, it is essential to musicologists to know the relationship to the *composers* which allows to include domain knowledge into the analysis process. Usually, one can assign musical pieces to a certain *type or style* which often correlates to a specific genre. Considering this contextual information can help detect and compare entities from a musical score collection. As discussed in section 2, there exist various frameworks that allow for extracting *statistical features* from symbolic music. We use music21 to extract jSymbolic features from digital sheet music [CA10] that provides the foundation for the visual analysis workspace that we introduce in the following sections.

Specifically, these low-level features contain rhythmic, melodic, and pitch information that provide informative attributes about single compositions or numerical distributions when investigating a score collection. In addition, we consider available contextual metadata including composer, composition type, and temporal characteristics.

Task Characteristics – Content-based MIR aims at uncovering music data relevant characteristics that can be employed for similarity computations and retrieval tasks [GN19]. For instance, a specific retrieval task **[T1]** would be the identification of all compositions from *Mozart* that have a time signature of a *triple meter*. Besides performing inquiring based on exact information, music analysts could also be interested in exploring similarities of pieces within a dataset **[T2]** with the objective to create new hypotheses or gain new insights about the underlying information [KKM*20]. Additionally, the comparison of two or even a set of compositions **[T3]** can be relevant during music analysis to identify commonalities of the complete work of composers or different versions of the same composition [ULMS10]. Not only comparison tasks can benefit from grouping and clustering based on meta information **[T4]**, but also the investigation of single composers based on statistical values, they provide a quick overview without the need to view all items separately. Depending on the quality of the dataset, it can be crucial for analysts to identify duplicates within a dataset **[T5]** to omit works from further analysis or other issues such as detecting musical plagiarism [Cam20]. Based on our cooperation with musicologists we focus on the following tasks that address typical analysis challenges:

[T1] Retrieve/Filter compositions based on title, composer, or type
[T2] Exploratory analysis of a musical score collection
[T3] Comparison of scores, composers, types, and features
[T4] Clustering/grouping of pieces, composers, and types
[T5] Detection of similar entities in a given corpus

Visualization Requirements – Based on the described data and task characteristics, we can identify essential requirements for the visualizations that are necessary to support the analysis process. For the information retrieval task **[T1]**, we need a separate filter and sort functionality for each meta feature and the low-level characteristics. Suppose a user is looking for a specific composer or composition. In that case, we need an interactive search bar that allows for direct querying of the underlying dataset to avoid the tedious search for data items within the provided visualizations. To associate the temporal aspects of composers, we need a timeline visualization that can also be directly used to filter composers to be analyzed by the user. The type of composition is another fundamental characteristic for music analysts that can be used for exploration **[T2]**, comparison **[T3]**, or identification of interesting pieces. To extract similar compositions within the dataset, we can employ a dimension reduction technique such as MDS [CC08] to project dozens of low-level features into a two-dimensional space **[T5]**. Then, identifying groups based on a clustering method is required to detect multiple pieces that have related features **[T4]**.

Used Dataset – MuseScore comprises a heterogeneous dataset with over a million compositions. Instead of using all these files, we manually selected a subset from certain users who upload curated files, e.g., ClassicMan [Cla21]. Typically, these compositions have high community ratings. Because of this preselection process, we ensure that the compositions in the corpus are of better quality.

4. Visual Interactive Analysis Workspace

This section introduces the interactive workspace and describes three exemplary analysis workflows to illustrate how to use it for analyzing sheet music corpora. CorpusVis is available as a web-based application via **https://visual-musicology.com/corpus**.

4.1. CorpusVis – User Interface

The analysis workspace consists of five components (see Figure 1) supporting data filtering and visual analysis following Shneiderman's Visual Information Seeking Mantra [Shn96]. The components are connected through animated linking and brushing, showing relevant details during the analysis process to the user. For instance, if the user hovers a circle (represents a composition or aggregations) in the projection view D, then the corresponding composer and forms A and matrix row(s) matrix B are highlighted. The settings widget (Figure 2, top left) enables users to switch between the German and English and provides a guided tour for all components.

Data Import and Selection – This component enables uploading custom sheet music datasets (*.mxl* format) or filtering existing datasets by keywords (composer names, title, etc.), composition types, and via composer selection from the timeline **[T1]**. The user can save an interesting data selection as a use case, or simply reset the analysis through further options provided in the settings
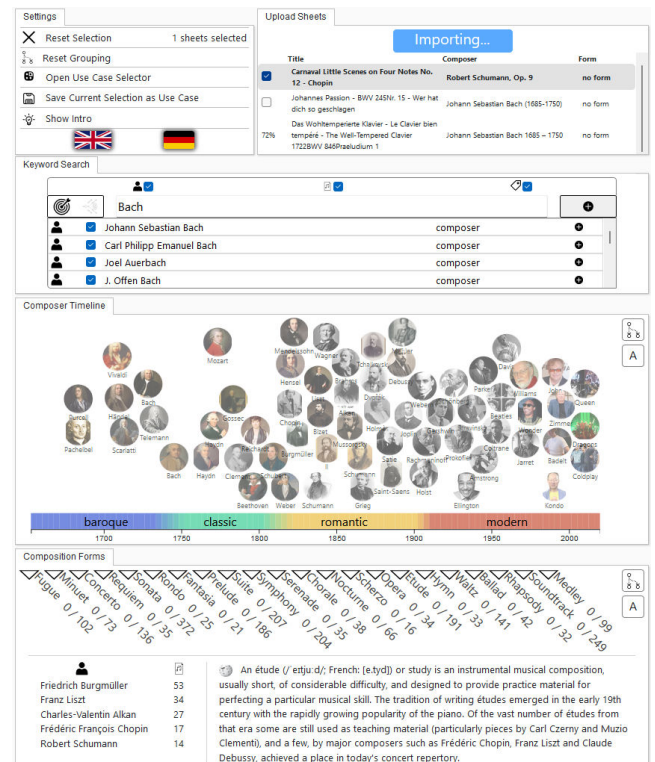

Figure 2: *The keyword search allows to filter composers and composition types. Users can extend the dataset by uploading custom data. Further settings: reset the system state, save use cases, or open an interactive introductory guide through through CorpusVis.*

Figure 3: *The feature distribution chart provides statistical characteristics (min., median, max.) of a subset and enables comparing the subset to the rest of the corpus regarding a selected feature.*

widget. The composer timeline shows 62 composers for which we could directly extract their life data from the files. Similar to Weiß et al. [WMDM18], we use four historical periods (baroque, classic, romantic, modern) to position the composers on the timeline. We apply an animated layout that uses the vertical axis to avoid overlapping.

jSymbolic Feature Matrix – Depending on the grouping, the feature matrix encodes statistical information extracted from single compositions or clusters at the workspace' center. User interaction allows sorting and filtering all features to focus only on subsets if needed **[T3]**. Analysts can also inform themselves by getting detailed information for each available feature. The features support the identification of salient patterns, and the matrix can be used for comparison tasks **[T2]**. Overall, the feature matrix can display up to 46 different jSymbolic features extracted with music21 [Cc21]. The feature distribution chart (see Figure 3) helps analysts compare the statistics of a feature to the whole corpus, including the minimum, median, and maximum value of the feature for each group **[T3]**. This example shows the statistical information of the feature *Note Density,* which corresponds to the average number of notes that are played within a second. The values show that the corpus contains pieces with a note density above 100, while the selection is situated has higher values at the left side of the distribution compared to the overall corpus. The distribution chart illustrates that the feature regarding the selection is similar to the rest of the corpus (left-skewed). Figure 4 displays 14 features horizontally sorted


Figure 4: *This exemplary matrix shows 14 features from the two composers Clara Schumann and Modest Petrovich Mussorgsky. Users can visually detect low (purple) and high (yellow) values.*
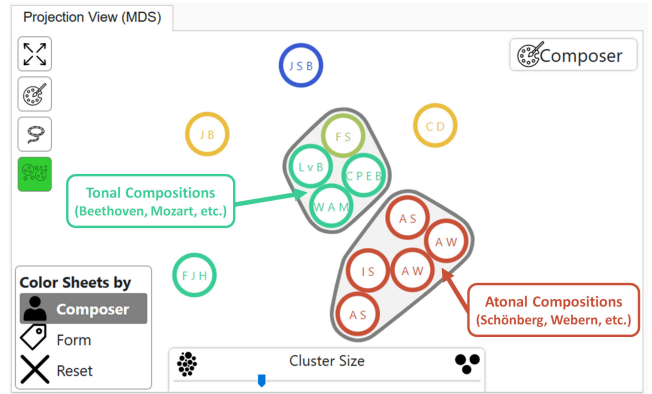

Figure 5: *MDS Projection View: If no grouping is applied, each circle in the projection view represents a single composition. The color provides the information to which epoch a piece is assigned.*

by category (M, P, R) with the rhythmic feature *Staccato Incidence* emphasizing a few pieces with a high amount of staccato usage.

MDS Projection View – We compute a multidimensional scaling projection based on the selection of low-level features to show the similarity between single compositions or clusters in a separate component **[T5]**. The circle labels displayed in Figure 5 are the initials of the composer names. In this case, the color of the circles provides epoch information, and the respective color scale is shown in the composer timeline. Analogously, a separate color scheme is used for the composition types. On-demand, the user can execute a DBSCAN clustering [HK98] and manipulate its *EPS* parameter via the cluster slide at the bottom to identify groups of similar entities **[T4]**. DB-SCAN has the advantages that it does not require a parameter for the number of clusters and it can identify outliers. We set MinPTS = 2 to enable the identification of duplicates. The visual grouping of the clusters uses concaveman [Aga17], a JavaScript library that builds on an algorithm based on work by Park and Oh [PO12] which allows to create concave hulls around given clusters.

Figure 6: *Analysts can use the metadata table to sort selected compositions and simply lookup interesting pieces from the list.*

Metadata Table – As a quick reference to the feature matrix, the workspace contains a meta information table of the current selection as displayed in Figure 6. Users can alphabetically sort a current selection based on composers, types, titles, and opus numbers. This rather simple data representation facilitates the connection to the rows of the more abstract feature matrix (see Figure 4) and projection view (see Figure 5).

Sheet View – Viewing sheet music using standard visual encoding is crucial for music analysts. Thus, the last component is the *sheet view,* which music analysts can use to estimate the quality of the underlying content (see Figure 7). Users can open the first page of compositions and relate it to the features via the context menu. Hence, the analyst can better understand abstract features by seeing the direct relation through the familiar notation. Via the playback icons at the top left corner of each item in the sheet view, analysts can also listen to the compositions on demand. We use the command-line version of MuseScore [Mus11] to convert the underlying MusicXML files into their .mp3 version. These files are also cached to reduce loading times if multiple users want to play back the same compositions.

4.2. Analysis Use Cases

CorpusVis offers multiple prepared analysis use cases consisting of preselected data and specific component configurations tailored to the analysis of the respective topic. Users can create additional scenarios by either manually filtering compositions from the prepared dataset compositions or uploading their own datasets. In the remainder of this section, we will introduce three exemplary use cases. The first two use cases represent a hypothetical user who is interacting with the system while `[UC3]` has been provided by one of our expert users, showcasing how CorpusVis covers the introduced tasks. All use cases can be opened directly in CorpusVis via the following links: Investigating Epoch Characteristics `[UC1]`, Tonality vs. Atonality `[UC2]`, and Composer Comparison `[UC3]`.

Investigation of Epoch Characteristics `[UC1]` – In this use case, the analyst wants to explore the differences between the historical epochs `[T2]`. Thus, from the list of use cases, the analyst selects *Epoch Comparison.* This use case considers 51 heterogeneous compositions from composers like *Pachelbel* (baroque), *Wagner* (romantic), and *Schönberg* (modern). Here, only a subset of eight features of the three feature categories (M, P, R) `B` such as *Repeated Notes* and *Pitch Variety* are selected that separate the four epochs within the projection view `[T3]`. While the modern and baroque compositions are distinguishable from the all others, the pieces from the intermediate epochs classic and romantic have more overlap `D`. The analyst identifies composers between the romantic and modern epoch, such as *Stravinsky*, *Rachmaninoff*, *Schönberg*, and *Webern*, that reside within classical and romantic composer groups. By using the column of the feature *Size of Melodic Arcs* to sort the feature matrix, the analyst detects that romantic composers (yellow circles in the projection view) often used large melodic arcs compared to artists from other epochs.

Atonality versus Tonality `[UC2]` – An interesting aspect regarding the harmonic development between the different epochs is the discussion of tonality and atonality. Atonal music is lacking a tonal center which means that none of the twelve available pitch classes is favored over the others as it is typically the case for classical music. Moreover, atonal music comprises more dissonant intervals (such as minor second and tritone intervals) which are available in the pitch feature group `B`. Specifically low values of the feature *Most Common Pitch Class Prevalence* are a good indicator for a composition to be atonal. In this example, the analyst is presented with 18 compositions from both classical composers including Mozart and Beethoven who composed rather tonal music as well as composers


Figure 7: *The sheet view allows viewing the first page of compositions. At the sheet level, analysts can perform harmonic and rhythmic analysis tasks using **MusicVis** [MFH*22].*

like Schönberg and Webern who are well-known by musicologists for their atonal compositions `A`. After starting with the analysis, the analyst uses the projection view `D` to see which compositions are positioned closer together to detect similar pieces `[T5]` (see Figure 5). As described, the analyst knows that tonality depends more on pitch and melodic features and not on rhythm. The configuration of the feature matrix `B` in this use case does not contain any rhythmic feature, which covers the mental model of the analyst. By viewing the columns in the feature matrix, the analyst detects as expected that the *Melodic Tritones* feature has much higher values for the atonal compositions confirming the hypothesis that atonal compositions contain more dissonant intervals `[T2]`. By investigating the feature distribution chart of *Number of Common Pitches*, the analyst discovers that the compositions subset of the configuration comprises more pieces with zero common pitches compared to the overall distribution in the corpus `[T3]`. Since the correlation matrix confirms a positive correlation of the *Number of Common Pitches* to *Most Common Pitch Class Prevalence*, the analyst is interested to see the values for the single compositions. The analyst is able to detect that the pieces by Schönberg and Webern try to avoid a tonal center as none of the available pitches seems to be favored compared to the works by the classical composers `[T5]`.

Composer Comparison `[UC3]` – In the first scenario, the domain expert compares similar pieces from two different composers `[T3]` and selects *Erlkönig* `[T1]` by Franz Liszt and Franz Schubert `A`. As a first step, the analyst manually inspects the first page using the sheet view `E` and detects the use of melodic octaves in the left hand of Liszt's interpretation which appears to be a salient difference between the compositions. Based on this observation, the expert focuses on the features *Melodic Octaves* and *Average Melodic Interval* which he assumes that these features should reflect this characteristic. The feature matrix `B` confirms the analyst's initial finding as the values for the *Melodic Octaves* feature is 28.1% for Liszt, but only 2.05% for Schubert `[T3]`. Consequently, the analyst wants to explore whether this aspect is a typical difference between the two composers `[T2]` and selects all available pieces in the dataset from both composers and aggregates them based on the composer attribute `[T4]`. Although the distinction decreases, the trend is still present (6.30% - Liszt vs. 3.8% - Schubert). While the

two features selected by the user are correlated, the analyst discovers that they represent discriminating characteristics between Liszt and Schubert. As the last step, the analyst is eager to see any unusual pieces (outliers) and ungroups the two clusters back into the single compositions. By activating the clustering in the MDS projection view **D** the analyst groups similar pieces into groups finding many similarities on the feature level between Liszt and Schubert **[T5]**. For instance, while the compositions *11. Ecossaise D.781* (Schubert) and *Vive Henri IV* (Liszt) are quite similar on the feature level. The user finds *Ave Maria - Ellens dritter Gesang* by Schubert to be an interesting outlier compared to most of the other compositions. Eventually, the analyst uses the metadata table **C** to sort the matrix using the composer attribute to explore differences between the two collections **[T2]**. The analyst confirms the less surprising hypothesis that Liszt's works have a much higher *Pitch Variety*. Based on the feature *Number of Common Pitches* he discovers that while Liszt almost always has zero pitches which individually account for at least 9% of all notes, Schubert even has compositions with five common pitches such as *Nachtviolen (D.752)* **[T3]**.

5. Pair Analytics Study

We conducted pair analytics sessions with users of different expertise levels to evaluate the applicability of CorpusVis introduced in section 4. Since the number of available qualified domain experts is limited, we decided to use the pair analytics study method [AKGF11].

Preliminary Expert Feedback – Before the actual study, we presented CorpusVis to three domain experts to gather initial feedback in an informal pre-study. The domain experts appreciated interacting with an exploratory visualization. One of the experts liked the way on how music information could be retrieved through simply entering keywords or mouse interactions on the composer timeline or types **[T1]**. The interlinked components helped in understanding the connection between low-level features and the compositions **[T2]**. They found it useful to explore similarities and differences between composers and types by grouping them together **[T3]** or through using the clustering in the projection view **[T4]**. They appreciated the multiple components and the use of color to visually show connections of similar entities **[T5]**. Especially, they could use the feature filtering functionality to identify outliers within single feature dimensions. They also noted that the detailed description of the features in connection with the sheet view provides music students the opportunity to experience new characteristics of sheet music **[T2]**. In their opinion, the sheet music component is a core element to verify hypotheses and to investigate the quality of compositions. However, they also critiqued our used color map and the lack of a feature to upload their own music sheets. Both of these issues were addressed in a refinement cycle before our study. In addition, based on our observations from the interactions in the pre-study, we added the features of interactive clustering and distribution-based grouping and filtering.

5.1. Study Design

We conducted the study to gather user feedback to identify benefits and drawbacks of the introduced analysis workspace. For the evaluation, we used a video conferencing software to show the interactive
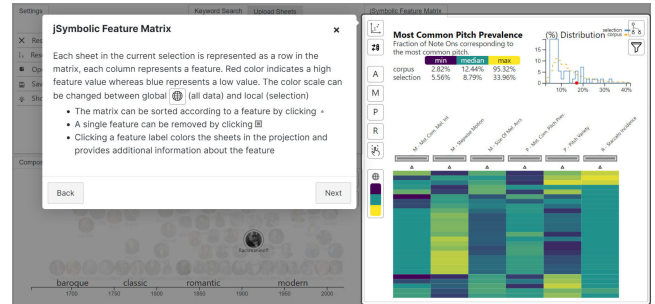

Figure 8: *We used Intro.js [Meh16] to prepare a visual explanation of all components of our prototype to ensure a consistent introduction of the interactive analysis workspace for all study participants.*

workspace to the participants during the study trials which enabled them to easily interact with CorpusVis ensuring equal study conditions. Based on the elicited *task characteristics* described in section 3, the participants carried out different analysis questions provided by the study director. From the study results, we gained useful insights on our work and its potential impact.

Study Participants – In sum, nine participants with different expertise levels partook in the study. Due to the required musical knowledge to use the analysis application, we focused our study on the target user group of domain experts including musicologists and music teachers. Five participants had an experience level of at least ten years including four university professors and a music teacher, which we further refer to as **s**ubject **m**atter **e**xperts (SME_1–SME_5). The remaining four participants were Ph.D. students (ST_1–ST_4) who are professionally involved with music on a daily basis but with less experience about music analysis compared to the SMEs.

Methodology – We applied a combined *observation* and *pair analytics* study [AKGF11] to investigate the usefulness of CorpusVis. In total, we performed nine study trials (average duration of 95 minutes) that were each led by a visual analytics expert (VAE) of our team. Each study session was structured into four sections.
The first part started with an *introduction* into the topic of sheet music collection analysis to frame the context of our work. Each participant completed a *demographic questionnaire* to provide information about their age, gender, level of education, and their level of expertise regarding music analysis. The first phase ended with an interview about the participants' *expectations* regarding an analysis system intended to support the investigation of sheet music collections.

After the general introduction, each participant received a detailed, textually guided *walkthrough* regarding all components of the study prototype (see Figure 8). Then, we introduced the five analysis tasks **[T1]**–**[T5]** from the analyzed requirements to each participant. Besides the analysis tasks, the introduction comprised an explanation of the relevant metadata features, including composer, type, epoch, and the low-level features extracted from the musical compositions visualized by the different components. In addition, the participants were informed about the origin of the used dataset (MuseScore [Mus11]) and its content addressing the overall amount of considered composers and musical pieces. Phase 2 closed with a simple and unguided exploration (similar to **[T2]**) of the visual interactive workspace, allowing the participants to get acquainted with the

application. This introductory task contained a preselected set of ten different compositions and five features. As a quick entry point, we prepared a use case *feature explanation* that can be directly loaded from the use case selector in the interface the provides this preselection. We chose the pieces based on contrasting (extreme) feature values such as low and high *Most Common Pitch Prevalence* values (3% and 65%). The participants were asked to explore the *Feature Matrix* (see Figure 4) to sort the selected sheets according to features of their interest. Then they used the *sheet view* (see Figure 7) to relate the feature values to the underlying common music notation. During the main part of the study, we asked the participants to execute the analysis tasks, as described in the next section. We encouraged them to verbalize their thoughts and interactions (*think-aloud concept*). After completing the prescribed analysis tasks, we performed an interview to receive feedback about the participants' experience regarding the functionality and usability of the prototype. We closed the study trial with a final questionnaire containing nine questions answered via a five-point Likert scale to identify the benefits and drawbacks of CorpusVis.

Study Tasks – During the pair analytics session, all participants analyzed sheet music collections based on three use cases introduced in subsection 4.2. [UC1] comprised 51 compositions from various composers of all epochs. After getting acquainted with the metadata of the selected pieces, the participants selected all pieces from a single epoch of choice [T1] and identified at least two discriminating features [T3]. For this task each participant used the feature matrix for identifying similarities between the compositions [T5]. In addition, the participants could add additional pieces through the keyword search to extend the selection [T1]. We encouraged each participant to freely explore the visualizations with individual compositions of interest [T2].

A second task [UC2] addressed the comparison of *atonal* with *tonal* pieces (Tonal versus Atonal), which used contemporary and classical works. Similarly, the participants were asked to detect characteristic features values for both groups by comparing the music sheets in the projection view and the feature matrix [T3]. The last task addressed comparing two or more composers relevant to the participant. For this, the *Composer Timeline* helped filtering the respective composers [T1], followed by a grouping of all pieces of single composers [T4]. Then, we asked the participant to visually examine the feature differences [T5] of the composers [T2].

5.2. Evaluation Results

We gathered feedback from the qualitative evaluation during the different study parts that we grouped thematically in the following.

Participants' Expectations – During the expectation interview, the participants stated various aspects of an application intended to support the musical analysis of sheet music corpora. ST_1 remarked that he expected the application to group a musical repertoire into genres and composers automatically. SME_1 expected the analysis system to use all information available in musical scores to perform musical analysis, which should present differences between musical styles, since "*classical is different from pop music*". SME_2 mentioned that she expects an interface that allows her to provide information about her interests. Then, the "*system would automatically provide analysis results without the need for manual analysis*".

Exploratory Analysis – To familiarize themselves with the compositions of [UC1], the participants started that by investigating the metadata table to see which composers are selected. Applying the color scheme of the composer timeline to the projection view revealed the similarity of the compositions. SME_1 mentioned the separate grouping of contemporary Jazz and Pop pieces compared to Baroque and Classical compositions, which also were separate groups but had some overlapping. Using the lasso to select the yellow-colored (modern) pieces, she identified that compared to the rest, these pieces have more note repetitions and high proportion of consonant melodic intervals and less minor/major second intervals. ST_4 checked the quality of a sheet file based on features such as pitch variety. Typically complex compositions such as etudes and symphonies have a high pitch variety. Thus, participants could delineate custom arrangements from more original compositions. While analyzing the configuration of [UC2], the participants were presented with tonal compositions from composers such as Mozart and Beethoven, which were compared with atonal pieces from Schönberg and Webern (see Figure 5). The participants used the projection view to compare similarities and appreciated the sorting functionality in the feature matrix, which helped them confirm that atonal pieces use tritone intervals more often. They identified that the feature *Number of Common Pitches* for the atonal compositions is always zero, which reflects that composers objective not to favor certain pitch classes. Especially SME_2 appreciated the sheet view, which allowed her to compare the values from the feature matrix to the common music notation. The third use case started with aggregating all compositions from the composers displayed in the composer timeline. The participants found that *Changes Of Meter* is the most discriminating feature revealing that contemporary music seldomly changes the meter signature within a composition.

Comparative Analysis – The participants used the feature matrix to identify similarities between single compositions or composer groups. SME_4 and SME_5 found the melodic features useful to investigate the melodic content. In addition, SME_3 and ST_1 appreciated the features regarding the intervals, which they considered to be particularly helpful to identify dominant melodic aspects. SME_1 mentioned that the matrix helped confirm that contemporary music such as "Imagine Dragons and Coldplay is more consonant and has more note repetitions, but fewer stepwise motion" compared to Schönberg (Op. 11) or Webern (Op. 4). Using CorpusVis, SME_1 learned that minuets from Bach have much lower *Pitch Class Variety* compared to his fugues. ST_1 noted that based on the available pieces, Schönberg's work has more chromatic motion (25%) compared to Webern (20%). ST_2 and ST_3 emphasized that the pitch range and density information helped them investigate the pitch material to better understand the complexity without even considering the sheet view. ST_3 could identify that Beethoven uses 20 fewer pitches compared to Liszt, who, as he stated, uses more chromatic melodic progressions. During the investigation, the participants found different matrix features to be relevant. They could easily understand the separation into the three categories rhythm, pitch, and melody, which were of different importance depending on the task. For instance, the time signature is a salient feature that helps to determine certain composition types. The pitch variety helped them detect complex compositions and differentiate tonal from atonal compositions. The chromatic motion feature supported identifying twelve-tone pieces.

Closing Questionnaire – All participants completed a final questionnaire answering nine questions to assess the usability of CorpusVis. The questions and the results that are based on a 5-point Likert scale ranging from "Strongly Disagree" to "Strongly Agree" are displayed in Figure 9. The questionnaire results reveal that our participants gained new insights (Q2) about sheet music collections by enabling them to perform groupings of the dataset to carry out comparison tasks based on epoch, composers, and composition type information (Q7). Most participants agreed that CorpusVis is suitable for the visual analysis of sheet music corpora (Q1, Q9). Q6 received the lowest scores, which illustrates that the computationally extracted features do not sufficiently address all users' aspects, such as dynamics, instrumentation, or harmonic progressions. Nevertheless, the scores of Q5 show that the computed statistical features still enable the comparison even though they may contain all aspects the analysts would be interested in. All but one of the participants stated that they would recommend the interactive workspace to a colleague or use it for their studies and research (Q4). For most participants, the currently used dataset in CorpusVis contained the compositions they wanted to add to their analysis (Q8). Overall, the questionnaire results convey that the participants found potential for analyzing sheet music collections and would continue using it in the future, especially if it is improved based on this study's outcomes.

6. Discussion

We identified benefits and drawbacks from the study which showed that CorpusVis supports specific analysis tasks such as comparing composition types or even the entire work of composers. Many participants questioned the quality of the underlying MuseScore dataset which contains volunteered sheet music. To address this issue, CorpusVis enables the upload of sheet music collections, allowing musicologists to use their own, curated datasets if desired. In this way, we mitigate the negative influence of datasets that suffer from poor quality. Despite the underlying dataset's limited quality, CorpusVis demonstrates the positive impact that tailored visualizations have on the work of musicologists.
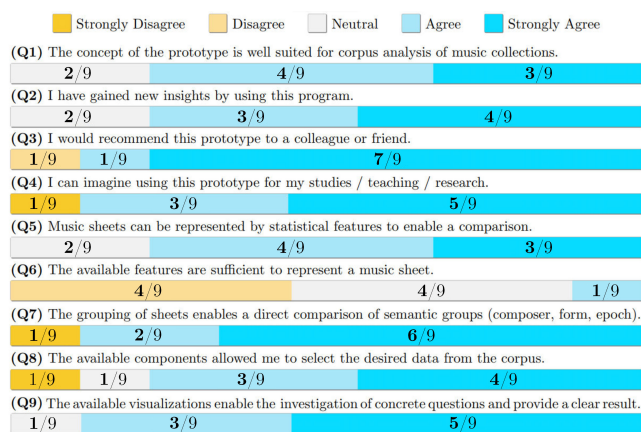

Figure 9: *The final questionnaire comprised nine questions to assess the usability and functionality of CorpusVis for analyzing sheet music collections. Most questions achieved good ratings, emphasizing an overall high satisfaction level.*

Musicologists controversially discuss whether and where to place clear borders between different epochs of music history. In our approach, we simplify the timeline of musical history into four major epochs *baroque*, *classic*, *romantic*, and *modern*. While we assign individual colors to each epoch, study participants appreciated that we introduce smooth color gradients to emphasize the transitions period. During the different musical periods, the focus on specific characteristics of music has shifted. While CorpusVis already includes several features, there may be other music styles that would require to extend the application. For instance, starting from the second half in the 20th century, new musical frontiers emerged, including serial, electronic, minimal, spectral, and experimental music. These contemporary music styles sometimes may not be available through transcribed common western music notation. They often rely on auditory or physical features such as timbre or spectral data. Therefore, such music compositions may only be available through audio files but not symbolic music. CorpusVis would either require an audio extension or a suitable symbolic notation that can be used to extract features that represent the underlying information.

During the implementation and design of the presented interactive workspace, we have encountered several challenges. A typical issue in this field is *copyright*. As MuseScore users sometimes upload compositions from contemporary musicians or artists who are still alive, these pieces may be deleted from MuseScore again. To avoid legal issues caused by sheet music that should not be visible in our system, we only include compositions from artists who died more than 70 years ago that are still available on MuseScore. Whenever sheets are deleted from MuseScore for any reason they are also no longer available in our system.

During the development of this system we were in close contact with multiple domain experts from musicology, allowing us to gather early feedback on CorpusVis. We quickly identified challenges in this inter-disciplinary collaboration, especially in creating visualizations that were understandable by experts while allowing them to focus on aspects relevant to their intended analysis. In particular, we found that the field of music analysis is extremely broad, necessitating a flexible system with a myriad of different options. However, adding additional functionality also complicated the system's design, in turn making it more difficult for experts to understand and use. As a result, in this work we focus on specific tasks and data characteristics to limit the scope.

6.1. Limitations

While the study results show that CorpusVis supports different music analysis tasks, it also surfaced some limitations. For instance, some composers of the 20th or 21st century heavily use timbral or textural elements. These features are often not encoded using the common Western music notation at all. As a result, CorpusVis cannot be of help in analyzing work from such contemporary composers. Future extensions of CorpusVis could integrate auditory features derived from recordings rather than only music sheets to capture the expression with which musicians played the compositions to close this gap. In its current implementation, CorpusVis does not capture the instrumentation of a composition. Several participants (SME_2–SME_5) stated that they would like to filter all compositions of a composer that contain piano or stringed instrument parts. Espe-

cially, identifying how the pitch range is distributed between different instruments would provide useful semantic information about the role of instruments and overall timbre of a piece.

While CorpusVis contains works from several composers, the visualization only partly scale based on the available aggregations. The projection view and feature matrix scale towards adding additional compositions from the same composer using aggregation based on the composers as it uses the equal visual space. Yet, an analyst may need to focus the analysis on a subset of musical features if other characteristics should be considered as the cognitive complexity increases with the number of features to be analyzed. Nevertheless, CorpusVis already uses semantic groups for melodic, rhythmic, and pitch features to enable users to focus on a subset. To improve this, CorpusVis could be adapted to allow users to create custom feature groups to focus the analysis and additional features that could be easily added if required. The application's performance depends on the library music21, which loads every composition once at the beginning to prepare all features for visualization.

SME_2 mentioned that the classification of the pieces into their types may be ambiguous, and that enabling analysts to re-classify compositions could increase the analyst's influence on a dataset's quality besides the option of uploading custom datasets. ST_1 and SME_1 remarked that the vertical combination of notes (simultaneously played notes is referred to as harmony) is not covered by the available features, since the employed jSymbolic features are limited to pitch, melodic, and rhythm characteristics. Although CorpusVis already offers a large array of selection and filtering mechanisms, some participants would have liked additional options for filtering a dataset by key signatures, voicing, and lyrics.

6.2. Take-Home Messages

From our work, we distill three primary take-home messages.

(1) Corpus-level music analysis is challenging due to the magnitude and complexity of underlying features. Manual analysis of sheet music collections is a tedious process requiring profound domain knowledge. CorpusVis supports the analysis of extensive sheet music collections, which is not feasible when done manually. While abstract visualizations open up new perspectives, they should not be considered a replacement for traditional analysis methods but rather an extension. Thus, integrating domain experts during the design phase and the requirements analysis is crucial to address their specific needs. This enables users to combine close- (sheet view) and distant-reading (projection, features) and get a corpus-overview.

(2) Visualization dashboards can enable multiple perspectives on the data, supporting scalable and multi-faceted analysis. Especially when musicologists want to explore, understand, and interpret music sheet corpora, visual analytics helps addressing the individual needs of analysts. Automatic computational methods do not provide visualizations to users facilitating interactive engagement with a sheet music collection to understand its underlying characteristics better. We allow for various interconnected analysis workflows by tailoring the components' visual design to their targeted task. The availability of multiple components covering different characteristics of the dataset enables users to set their analysis focus on those attributes they are interested in their analysis.

(3) Tailored interactions enable a task-driven analysis of corpora using diverse metadata and statistical features. Interactive visualization allows investigating sheet music from different perspectives, supporting various aspects music analysts may be interested in. In particular, for analyzing complex questions involving multiple meta-data and features, we rely on the interactive interlinking of visual components, enabling an expressive yet visually uncomplicated analytics system. Enabling analysts to import their own curated datasets facilitates their investigation and generation and confirmation of hypotheses.

6.3. Future Work and Research Opportunities

We postulate that our work only represents a starting point for future cooperation between visualization and musicology researchers, and we want to emphasize that it is worth "risking the drift" for such interdisciplinary collaborations [HEAB*17]. So far, the classification of the compositions metadata available in CorpusVis is fixed and cannot be adjusted by the analyst. In the future, it may be worthwhile to enable users to change meta information to increase the dataset quality. Extending the application through labeling would further improve the analysts' ability to capture and externalize analysis results. While we have considered the temporal aspects of music, we have neglected the geographic information of where composers have lived and worked which has been already studied by Khulusi et al. [KKFJ20]. It would be interesting to investigate how the location influenced the style of specific composers. While the low-level characteristics are just statistical measures, it would be interesting to see how they correlate with the mood and emotions of listeners. So far, we have not considered the dynamics, instrumentation, or timbre of the instruments that the compositions are written for, which would be worthwhile to add support for in the future. If pianists or guitarists are looking for new pieces to play, it would be helpful to provide specific instrument filters that help them find relevant information. Moreover, as the dataset's quality plays a primary role in analysis, it would be helpful to provide classification means for analysts to identify low-quality pieces to exclude unwanted items. Eventually, CorpusVis focuses the analysis on a high level. We recently published MusicVis that enables analysts to investigate single compositions at the sheet level [MFH*22]. We aim at connecting both approaches to allow analysts to filter relevant compositions at a higher level using CorpusVis which can then be further analyzed at the sheet level.

7. Conclusion

In this paper, we have introduced CorpusVis, a web-based interactive workspace for the visual analysis of sheet music collections. We discussed data and task characteristics relevant to analyzing sheet music corpora. Multiple connected components through linking and brushing support exploration, detection, and comparison tasks. Based on initial expert feedback, we identified benefits and challenges of our approach. In addition to three exemplary use cases that demonstrate the analysis workflow to compare composers and types, we conducted a pair of analytics study with nine participants. The qualitative study results show that music analysts benefit from interactive visualization addressing different use case scenarios such as comparing composers, epochs, and composition types based on low-level features including pitch, rhythm, and melody.

References

[Aga17] AGAFONKIN V.: A very fast 2D concave hull algorithm in JavaScript, 2017. Retrieved March 9, 2022, from `https://github.com/mapbox/concaveman`. 5

[AKGF11] ARIAS-HERNÁNDEZ R., KAASTRA L. T., GREEN T. M., FISHER B. D.: Pair Analytics: Capturing Reasoning Processes in Collaborative Visual Analytics. In *Conf. on Syst. Sc.* (2011), pp. 1–10. `doi:10.1109/HICSS.2011.339`. 7

[ANC17] ARMENTANO M. G., NONI W. A. D., CARDOSO H. F.: Genre classification of symbolic pieces of music. *J. Intell. Inf. Syst. 48*, 3 (2017), 579–599. `doi:10.1007/s10844-016-0430-7`. 2

[BBBD20] BLUDAU M., BRÜGGEMANN V., BUSCH A., DÖRK M.: Reading Traces: Scalable Exploration in Elastic Visualizations of Cultural Heritage Data. *Comput. Graph. Forum 39*, 3 (2020), 77–87. 3

[BEAC*18] BRADLEY A. J., EL-ASSADY M., COLES K., ALEXANDER E., CHEN M., COLLINS C., JÄNICKE S., WRISLEY D. J.: Visualization and the Digital Humanities. *IEEE Comput. Graph. and Appl. 38*, 6 (2018), 26–38. 2

[BHF*13] BOGDANOV D., HARO M., FUHRMANN F., XAMBÓ A., GÓMEZ E., HERRERA P.: Semantic audio content-based music recommendation and visualization based on user preference examples. *Inf. Process. Manag. 49*, 1 (2013), 13–33. `doi:10.1016/j.ipm.2012.06.004`. 2

[BMHC16] BRADLEY A. J., MEHTA H., HANCOCK M., COLLINS C.: Visualization, Digital Humanities, and the Problem of Instrumentalism. *Proc. of the VIS4DH at IEEE VIS 2016* (2016). 3

[CA10] CUTHBERT M. S., ARIZA C.: Music21: A Toolkit for Computer-Aided Musicology and Symbolic Music Data. In *Proc. of ISMIR* (2010), Int. Soc. for Music Information Retrieval, pp. 637–642. 3

[Cam20] CAMERON S.: *What Is Plagiarism and What Is Musical Plagiarism?* Springer, 2020, pp. 1–36. `doi:10.1007/978-3-030-42109-0_1`. 4

[CB09] CHEN Y., BUTZ A.: MusicSim: Integrating Audio Analysis and User Feedback in an Interactive Music Browsing UI. In *Proc. of the Int. Conf. on Intelligent User Interfaces* (2009), ACM, pp. 429–434. 3

[CC08] COX M. A., COX T. F.: Multidimensional Scaling. In *Handbook of Data Visualization*. Springer, 2008, pp. 315–347. 4

[Cc21] CUTHBERT M. S., CUTHBERTLAB: music21.features.jSymbolic, 2021. Retrieved July 8, 2021, from `https://web.mit.edu/music21/doc/moduleReference/moduleFeaturesJSymbolic.html`. 3, 5

[Cla21] CLASSICMAN: Sheet Music Collection of MuseScore user ClassicMan, 2021. Retrieved July 29, 2021, from `https://musescore.com/user/19710/sheetmusic`. 4

[CR16] CORRÊA D. C., RODRIGUES F. A.: A survey on symbolic data-based music genre classification. *Expert Syst. Appl. 60* (2016), 190–210. `doi:10.1016/j.eswa.2016.04.008`. 3

[FLTZ11] FU Z., LU G., TING K. M., ZHANG D.: A Survey of Audio-Based Music Classification and Annotation. *IEEE Trans. Multim. 13*, 2 (2011), 303–319. `doi:10.1109/TMM.2010.2098858`. 2

[GMB*14] GIPP B., MEUSCHKE N., BREITINGER C., PITMAN J., NÜRNBERGER A.: Web-based Demonstration of Semantic Similarity Detection Using Citation Pattern Visualization for a Cross Language Plagiarism Case. In *Proc. of the Int. Conf. on Enterprise Information Systems* (2014), vol. 2, SciTePress, pp. 677–683. `doi:10.5220/0004985406770683`. 3

[GN19] GEORGES P., NGUYEN N.: Visualizing music similarity: clustering and mapping 500 classical music composers. *Scientometrics 120*, 3 (2019), 975–1003. `doi:10.1007/s11192-019-03166-0`. 3, 4

[Goo01] GOOD M.: MusicXML for Notation and Analysis. *The Virtual Score: Representation, Retrieval, Restoration 12*, 113-124 (2001), 160. 2, 3

[HEAB*17] HINRICHS U., EL-ASSADY M., BRADLEY A. J., FORLINI S., COLLINS C.: Risk the Drift! Stretching Disciplinary Boundaries through Critical Collaborations between the Humanities and Visualization. In *Proc. of the IEEE VIS Workshop on Vis. for the Digital Humanities* (2017). 10

[HK98] HINNEBURG A., KEIM D. A.: An Efficient Approach to Clustering in Large Multimedia Databases with Noise. In *Proc. of the Int. Conf. on Knowledge Discovery and Data Mining* (1998), AAAI Press, pp. 58–65. URL: `http://www.aaai.org/Library/KDD/1998/kdd98-009.php`. 5

[JFCS15] JÄNICKE S., FRANZINI G., CHEEMA M. F., SCHEUERMANN G.: On Close and Distant Reading in Digital Humanities: A Survey and Future Challenges. In *Eurographics Conf. on Vis.* (2015), pp. 83–103. 2

[JFCS17] JÄNICKE S., FRANZINI G., CHEEMA M. F., SCHEUERMANN G.: Visual Text Analysis in Digital Humanities. *Comput. Graph. Forum 36*, 6 (2017), 226–250. `doi:10.1111/cgf.12873`. 3

[JFS16] JÄNICKE S., FOCHT J., SCHEUERMANN G.: Interactive Visual Profiling of Musicians. *IEEE Trans. on Vis. and Comp. Graph. 22*, 1 (2016), 200–209. `doi:10.1109/TVCG.2015.2467620`. 3

[Kir07] KIRSCHENBAUM M. G.: The remaking of reading: Data mining and the digital humanities. In *The National Sci. Found. Symp. on Next Gen. of Data Mining and Cyber-Enabled Disc. for Innovation* (2007), vol. 134. 3

[KKFJ20] KHULUSI R., KUSNICK J., FOCHT J., JÄNICKE S.: musiXplora: Visual Analysis of a Musicological Encyclopedia. In *Proc. of the Int. Joint Conf. on Comput. Vision, Imag. and Comput. Graph. Theory and Appl.* (2020), SciTePress, pp. 76–87. `doi:10.5220/0008977100760087`. 3, 10

[KKM*20] KHULUSI R., KUSNICK J., MEINECKE C., GILLMANN C., FOCHT J., JÄNICKE S.: A Survey on Visualizations for Musical Data. *Comput. Graph. Forum* (2020). `doi:10.1111/cgf.13905`. 4

[KO07] KEIM D. A., OELKE D.: Literature Fingerprinting: A New Method for Visual Literary Analysis. IEEE Comput. Soc., pp. 115–122. 3

[LME*11] LIEM C. C., MÜLLER M., ECK D., TZANETAKIS G., HANJALIC A.: The Need for Music Information Retrieval with User-Centered and Multimodal Strategies. In *Proc. of the Int. ACM Workshop on Music Information Retrieval with User-Centered and Multimodal Strategies* (2011), ACM Press. `doi:10.1145/2072529.2072531`. 3

[LWSF20] LENAIN R., WESTON J., SHIVKUMAR A., FRISTED E.: Surfboard: Audio Feature Extraction for Modern Machine Learning. In *Interspeech, Conf. of the Int. Speech Comm. Assoc.* (2020), ISCA, pp. 2917–2921. `doi:10.21437/Interspeech.2020-2879`. 3

[MCF18] MCKAY C., CUMMING J., FUJINAGA I.: JSYMBOLIC 2.2: Extracting features from symbolic music for use in musicological and MIR research. In *Proc. of the Int. Soc. of Music Inf. Retrieval Conf.* (2018), Gómez E., Hu X., Humphrey E., Benetos E., (Eds.), pp. 348–354. URL: `http://ismir2018.ircam.fr/doc/pdfs/26_Paper.pdf`. 2, 3

[Meh16] MEHRABANI A.: Intro.js: Better Introductions for Websites and Features with a Step-by-Step Guide for Your Projects., 2016. Retrieved November 29, 2021, from `https://introjs.com/`. 7

[MF09] MCKAY C., FUJINAGA I.: jMIR: Tools for Automatic Music Classification. In *Proc. of the Int. Comp. Music Conf.* (2009), Mich. Publ. URL: `http://hdl.handle.net/2027/spo.bbp2372.2009.014`. 3

[MFH*22] MILLER M., FÜRST D., HAUPTMANN H., KEIM D. A., EL-ASSADY M.: Augmenting Digital Sheet Music through Visual Analytics. *Computer Graphics Forum* (1 2022). `doi:10.1111/cgf.14436`. 6, 10

[MGD11] MÜLLER M., GOTO M., DIXON S.: Multimodal Music Processing (Dagstuhl Seminar 11041). *Dagstuhl Reports 1*, 1 (2011), 68–101. 3

[MLCM16] MCCURDY N., LEIN J., COLES K., MEYER M. D.: Poemage: Visualizing the Sonic Topology of a Poem. *IEEE Trans. on*

Vis. and Comp. Graph. 22, 1 (2016), 439–448. `doi:10.1109/TVCG.2015.2467811`. 3

[MSK*19]   MILLER M., SCHÄFER H., KRAUS M., LEMAN M., KEIM D. A., EL-ASSADY M.: Framing Visual Musicology through Methodology Transfer. *Proc. of the IEEE VIS Workshop on VIS4DH* (2019). 2, 3

[Mus11]   MUSESCORE: MuseScore.com, 2011. Retrieved June 29, 2021, from `https://musescore.com/`. 2, 6, 7

[NN03]   NIENHUYS H.-W., NIEUWENHUIZEN J.: LilyPond, a System for Automated Music Engraving. In *Proc. of the Coll. on Musical Informatics* (2003), vol. 1, Citeseer, pp. 167–171. 2

[OKT21]   OKTAV: OKTAV, 2021. Retrieved July 8, 2021, from `https://www.oktav.com/`. 2, 3

[Pam01]   PAMPALK E.: *Islands of Music Analysis, Organization, and Visualization of Music Archives*. CiteSeer, 2001. 3

[PMK10]   PAULUS J., MÜLLER M., KLAPURI A.: State of the Art Report: Audio-Based Music Structure Analysis. In *Proc. of the Int. Soc. of Music Inf. Retrieval Conf.* (2010), pp. 625–636. URL: `http://ismir2010.ismir.net/proceedings/ismir2010-107.pdf`. 2

[PO12]   PARK J., OH S.: A New Concave Hull Algorithm and Concaveness Measure for n-dimensional Datasets. *J. Inf. Sci. Eng. 28*, 3 (2012), 587–600. URL: `http://www.iis.sinica.edu.tw/page/jise/2012/201205_10.html`. 5

[Sap05]   SAPP C. S.: Online Database of Scores in the Humdrum File Format. In *Proc. of the Int. Soc. of Music Inf. Retrieval Conf.* (2005), pp. 664–665. URL: `http://ismir2005.ismir.net/proceedings/3123.pdf`. 2

[Sap15]   SAPP C. S.: Graphic to Symbolic Representations of Musical Notation. In *Proc. of the Int. Conf. on Technologies for Music Notation and Representation* (2015), pp. 124–132. 2

[Sch19]   SCHEDL M.: Deep Learning in Music Recommendation Systems. *Frontiers in Applied Mathematics and Statistics 5* (2019), 44. 3

[SF20]   SHATRI E., FAZEKAS G.: Optical Music Recognition: State of the Art and Major Challenges. In *Proc. of the Int. Conf. on Technologies for Music Notation and Representation* (2020), Hamburg University for Music and Theater, pp. 175–184. 2

[Shn96]   SHNEIDERMAN B.: The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In *Proc. of the IEEE Symp. on Visual Languages* (1996), IEEE Computer Society, pp. 336–343. 4

[SSS*14]   SACHA D., STOFFEL A., STOFFEL F., KWON B. C., ELLIS G. P., KEIM D. A.: Knowledge Generation Model for Visual Analytics. *IEEE Trans. on Vis. and Comp. Graph. 20*, 12 (2014), 1604–1613. `doi:10.1109/TVCG.2014.2346481`. 2, 3

[STOH21]   SAKURAI K., TOGO R., OGAWA T., HASEYAMA M.: Music Playlist Generation Based on Graph Exploration Using Reinforcement Learning. In *IEEE Global Conf. on Life Sciences and Technologies* (2021), IEEE, pp. 53–54. `doi:10.1109/LifeTech52111.2021.9391870`. 2

[ULMS10]   URBANO J., LLORÉNS J., MORATO J., SÁNCHEZ-CUADRADO S.: Melodic Similarity through Shape Similarity. In *Exploring Music Contents - Int. Symp., CMMR 2010* (2010), Ystad S., Aramaki M., Kronland-Martinet R., Jensen K., (Eds.), vol. 6684 of *Lecture Notes in Computer Science*, Springer, pp. 338–355. `doi:10.1007/978-3-642-23126-1\_21`. 4

[VT19]   VERMA H., THICKSTUN J.: Convolutional composer classification. In *Proc. of the Int. Soc. of Music Inf. Retrieval Conf.* (2019), Flexer A., Peeters G., Urbano J., Volk A., (Eds.), pp. 549–556. URL: `http://archives.ismir.net/ismir2019/paper/000066.pdf`. 3

[WJVR16]   WEIHS C., JANNACH D., VATOLKIN I., RUDOLPH G.: *Music Data Analysis: Foundations and Applications*. Chapman and Hall/CRC Press, 2016. 2

[WMDM18]   WEISS C., MAUCH M., DIXON S., MÜLLER M.: Investigating style evolution of Western classical music: A computational approach. *Musicae Scientiae 23* (2018), 486–507. `doi:10.1177/1029864918757595`. 3, 5