Analysis of Timbral Augmentation in the Orchestration Analysis & Research Database (OrchARD)

Corinne Darche

Music Technology Area
Department of Music Research
Schulich School of Music
McGill University
Montréal, Canada

August 15, 2023

A thesis submitted to McGill University in partial fulfillment of the requirements for the degree of Master of Arts.

Abstract

Instrumental blend is a subject of great interest in both orchestration treatises and music perception research. Music psychologists have coined terms to explain the perceptual process of grouping different instrumental timbres together to create musical form. Although research has identified which instruments blend effectively, there remains much work to be done in applying these auditory grouping processes and blend theories to larger musical works. To promote the increased use of auditory grouping processes in music analyses, it is of interest to investigate potential instrumentation patterns in certain orchestral effects. As such, the current study presents a corpus analysis of the Orchestration Analysis & Research Database (OrchARD) to investigate instrumentation patterns in timbral augmentation blends. In timbral augmentation certain instruments are more likely to dominate a blend perceptually or to augment or embellish the dominant timbre. In timbral emergence, all instruments are subsumed into the emergent timbre. The aim was to explore the conditions which contribute to these different role assignments. Three different analysis techniques were conducted on the database's timbral augmentation entries, and occasionally the timbral emergence entries, to observe the differences. First, the relative frequencies of occurrence for each blend role (i.e., dominating instrument, embellishing instrument, emerging instrument) were calculated to determine initial patterns. Then, using the frequent pattern mining algorithm FP-Growth, these patterns were expanded on to find larger common groupings and cases of statistical dependence. The influence of the blend's strength, as rated by human annotators, was introduced at this step as well. Finally, a network analysis was conducted solely on the timbral augmentation data to explore the instrument relationships and potential subgroupings within the blend roles. The results from the analyses complement each other and share similarities with what is written in late 19$^{\text{th}}$-century and 20$^{\text{th}}$-century orchestration treatises. Certain instrument families were identified as more likely candidates for certain blend roles. For example, strings are more likely to be used as dominating instruments in timbral augmentation blends, and woodwinds are more likely to be used as embellishing instruments. Certain frequent instrument groupings, such as the oboe, flute, and clarinet combination for embellishing, were highlighted in orchestration treatises. Although these results are restricted by the limitations posed by the available data in

OrchARD, they provide a statistical foundation to what was put forward by these orchestration treatises and indicate potential instrumentation patterns contributing to the perceived timbral augmentation blend effect.

Résumé

Le mélange instrumental est un sujet de grand intérêt dans les traités d'orchestration ainsi que la recherche sur la perception musicale. Les psychologues de la musique ont créé des termes afin d'expliquer le processus perceptif du groupement des différents timbres instrumentaux pour créer une forme musicale. Bien que la recherche ait identifié quels instruments mélange effectivement, il reste beaucoup de travail à faire dans l'application de ces processus de groupement auditif et ces théories de mélange aux grandes œuvres musicales. Pour encourager l'utilisation accrue des processus de groupement auditif dans l'analyse musicale, il est important d'étudier les motifs potentiels de l'instrumentation dans certains effets orchestraux. Ainsi, l'étude en cours présente une analyse du corpus OrchARD (*Orchestration Analysis & Research Database*) afin d'enquêter les motifs de l'instrumentation des augmentations timbrales. Dans l'augmentation timbrale, certains instruments sont plus susceptibles de dominer perceptivement un mélange ou d'augmenter ou d'embellir le timbre dominant. Dans l'émergence timbrale, tous les instruments sont intégrés dans le timbre émergent. Le but était d'explorer les conditions qui contribuent à l'attribution ces rôles différents. Trois techniques d'analyse étaient menées en utilisant les entrées d'augmentation timbrale dans la base de données et parfois les entrées d'émergence timbrale afin d'observer les différences. Pour commencer, les fréquences relatives de présence pour chaque rôle de mélange (c.-à-d., instrument dominant, instrument embellissant, instrument émergent) étaient calculées dans le but d'identifier des motifs initiaux. Ensuite, grâce à l'algorithme d'extraction de motifs fréquents *FP-Growth*, les modèles étaient développés davantage pour trouver des plus grands groupements communs d'instruments ainsi que des cas de dépendance statistique. L'influence de la puissance du mélange, telle que jugée par les annotateurs humains, était introduite à cette étape aussi. Finalement, une analyse de réseaux était menée seulement avec les données d'augmentation timbrale afin d'examiner les relations entre les instruments et les sous-groupements potentiels dans les rôles de mélange. Les résultats des analyses se complètent et partagent des semblances avec les textes des traités d'orchestration de la fin du 19$^{\text{e}}$ siècle et du 20$^{\text{ième}}$ siècle. Certaines familles d'instruments avaient tendance à adopter certains rôles de mélange. Par exemple, les cordes étaient plus susceptibles à être utilisés comme des instruments dominants et les bois étaient plus

souvent utilisés comme des instruments embellissants. De plus, certains groupements fréquents d'instruments, comme la combinaison de hautbois, de la flûte traversière et de la clarinette pour embellir, étaient souligner dans les traités d'orchestration. Malgré les limitations des résultats posées par les limites des données disponibles dans OrchARD, ils fournissent une base statistique pour ce que proposaient les traités d'orchestration et indiquent les motifs d'instrumentation potentiels qui contribuent à l'effet de mélange du type augmentation timbrale.

Acknowledgements

This thesis would not have been possible without the support of my colleagues and loved ones. To Stephen McAdams: thank you for your guidance from the very start of my graduate studies. I entered Music Technology as a computer science student who was eager to learn more about timbre. You were gracious enough to open the door for me and make it a possibility. Words will never be able to fully express my gratitude for that. To Ichiro Fujinaga: thank you for your guidance throughout my thesis research and writing process, even when it meant navigating the large time difference between Montreal and Japan. You brought out the best in my work, and I cannot thank you enough for that. To Bennett Smith: thank you for helping me navigate the maze that is the MPCL server and for the many philosophical conversations shared in the lab. To my MPCL colleagues (Yifan, Linglan, Joshua, Jade, Erica, Yuval, Iza, Lena, Behrad, Félix, Ben, and Andrés): thank you for making the MPCL feel like a true family. Every day, I walked into the lab knowing that I'd be entering a supportive and welcoming environment filled with other timbre geeks, and I felt like the luckiest person in the world. Our dynamic is rare, and I will forever cherish my time as an MPCL-er. To the *eZone* (Paul, Brady, Margaret): thank you for keeping up the morale during coursework and for carrying the grad school stress with me. I cannot wait to see what exciting journeys you all take in the years to come. To my friends outside the four walls of the Music Technology area: thank you for being my moral support throughout my time at McGill. I cannot imagine where I would be without all of you. I'm grateful for the solidarity, the laughs, and the intense board game nights. To my family: thank you for your unending support. You trusted that I would figure out how to get there somehow, even when I didn't trust myself.

Author Contributions

Under the supervision of Professor Stephen McAdams and Professor Ichiro Fujinaga, I was responsible for coming up with the research questions, extracting the data from OrchARD, designing and running each data analysis technique, and analyzing and interpreting the results.

Contents

Abstract ii

Résumé iv

Acknowledgements v

Author Contributions vi

List of Figures xi

List of Tables xiii

1 Introduction 1

1.1 Perceptual Theories of Orchestration 2

1.1.1 Instrumental Blend 4

1.2 Previous Work: Instrumentation Corpus Studies 6

1.3 The Current Study 8

1.3.1 Objectives 8

1.3.2 The Orchestration Analysis & Research Database (OrchARD) 10

1.3.3 Overview 14

2 Experiment 1: The Relative Frequencies of Occurrence of Different Types of Blend 17

2.1 Method 18

2.1.1 Basic statistics of dataset 18

2.1.2 Procedure 22

2.2 Results 24

2.2.1 Analysis of blend roles 24

2.2.2 Analysis of case-study instruments 35

2.3 Discussion . . . 44

3 Experiment 2: Frequent Pattern Mining Using FP-Growth 49

3.1 Background on frequent pattern mining . . . 50

3.1.1 Algorithms . . . 51

3.2 Method . . . 54

3.2.1 Algorithm Implementation . . . 55

3.3 Results . . . 56

3.3.1 Analysis of blend roles . . . 56

3.3.2 Analysis of case-study instruments . . . 59

3.4 Discussion . . . 66

4 Experiment 3: Network Analysis of Timbral Augmentation Blends 71

4.1 Background on network science . . . 72

4.1.1 Terminology . . . 72

4.1.2 Centrality . . . 74

4.1.3 Community detection . . . 75

4.2 Method for network construction and analysis . . . 77

4.2.1 Network design . . . 78

4.2.2 Procedure for network construction and analysis . . . 79

4.3 Network analysis results . . . 80

4.3.1 Analysis of blend roles . . . 81

4.3.2 Analysis of case studies . . . 94

4.4 Discussion . . . 113

5 Conclusion 117

5.1 General discussion . . . 117

5.2 Limitations . . . 121

5.3 Future directions . . . 122

References 125

List of Figures

1.1 Diagram of auditory grouping processes and their related perceptual qualities and orchestral effects. Taken from McAdams et al. (2022) 3
1.2 Example of OrchARD annotation as seen on the current OrchARD user interface. The different colours on the score's page reflect different types of orchestral effects found in the passage. The timbral augmentation shown here is represented by the orange box on the bottom half of this page. 11

2.1 Blend Strength Distribution for Timbral Augmentation Blends 19
2.2 Blend Strength Distribution for Timbral Emergence Blends 20
2.3 Distribution of Instrument Families in Dominating Instruments 24
2.4 Distribution of Woodwind Instruments in Dominating Instruments 25
2.5 Distribution of Brass Instruments in Dominating Instruments 25
2.6 Distribution of Percussion Instruments in Dominating Instruments 26
2.7 Distribution of String Instruments in Dominating Instruments 27
2.8 Distribution of Instrument Families in Embellishing Instruments 28
2.9 Distribution of Woodwind Instruments in Embellishing Instruments 28
2.10 Distribution of Brass Instruments in Embellishing Instruments 29
2.11 Distribution of Percussion Instruments in Embellishing Instruments 29
2.12 Distribution of String Instruments in Embellishing Instruments 30
2.13 Distribution of Instrument Families in Timbral Emergence Blends 32
2.14 Distribution of Woodwind Instruments in Emerging Instruments 33
2.15 Distribution of Brass Instruments in Emerging Instruments 33
2.16 Distribution of Percussion Instruments in Emerging Instruments 34
2.17 Distribution of String Instruments in Emerging Instruments 34
2.18 Distribution of Instruments Playing with the Violin in Timbral Augmentation Blends 36
2.19 Distribution of Instruments Playing with the Violin in Timbral Emergence Blends 36
2.20 Distribution of Instruments Playing with the Flute in Timbral Augmentation Blends 38

2.21 Distribution of Instruments Playing with the Flute in Timbral Emergence Blends . 38
2.22 Distribution of Instruments Playing with the Oboe in Timbral Augmentation Blends 40
2.23 Distribution of Instruments Playing with the Oboe in Timbral Emergence Blends . 41
2.24 Distribution of Instruments Playing with the with French Horn in Timbral Augmentation Blends 42
2.25 Distribution of Instruments Playing with the French Horn in Timbral Emergence Blends 43

3.1 The Apriori algorithm pseudocode. Taken from Hegland (2007) 51
3.2 Diagram of an example dataset and its FP-Tree representation. Taken from Grahne and Zhu (2003) 53
3.3 Pseudocode for the FP-Growth algorithm, taken from Aggarwal (2015) 54

4.1 An example of an undirected network (left) and a directed network (DiGraph) (right). Taken from Barabási (2016) 72
4.2 An example of a bipartite network. Taken from Barabási (2016) 72
4.3 The difference between the different centrality algorithms. Assume that all edges have a weight of 1. Taken from Needham and Hodler (2019) 74
4.4 A simple undirected network with three groups 76
4.5 Dominating Instruments Network with Communities Formed by the CNM Algorithm and the Louvain Algorithm 83
4.6 Embellishing Instruments Network with Communities Formed by the CNM Algorithm 86
4.7 Embellishing Instruments Network with Communities Formed by the Louvain Algorithm 87
4.8 Timbral Augmentation Network with Communities Formed by the CNM Algorithm 91
4.9 Timbral Augmentation Network with Communities Formed by the Louvain Algorithm 92
4.10 Simplified Timbral Augmentation Network using only Instrument Family Nodes . . 93
4.11 Violin Dominating Instruments Network with Communities Formed by the CNM Algorithm and the Louvain Algorithm 95
4.12 Violin Embellishing Instruments Network with Communities Formed by the CNM Algorithm 98
4.13 Violin Embellishing Instruments Network with Communities Formed by the Louvain Algorithm 99
4.14 Flute Dominating Instruments Network with Communities Formed by the CNM Algorithm and the Louvain Algorithm 101

4.15 Flute Embellishing Instruments Network with Communities Formed by the CNM Algorithm . . . 103
4.16 Flute Embellishing Instruments Network with Communities Formed by the Louvain Algorithm . . . 103
4.17 Oboe Dominating Instruments Network with Communities Formed by the CNM Algorithm and the Louvain Algorithm . . . 105
4.18 Oboe Embellishing Instruments Network with Communities Formed by the CNM Algorithm . . . 107
4.19 Oboe Embellishing Instruments Network with Communities Formed by the Louvain Algorithm . . . 108
4.20 French Horn Dominating Instruments Network with Communities Formed by the CNM Algorithm and the Louvain Algorithm . . . 110
4.21 French Horn Embellishing Instruments Network with Communities Formed by the CNM Algorithm . . . 112
4.22 French Horn Embellishing Instruments Network with Communities Formed by the Louvain Algorithm . . . 112

List of Tables

1.1 Number of blend entries in OrchARD by composer 12
1.2 Pieces and movements featured in OrchARD's timbral augmentation and timbral emergence blend data 13

2.1 Distribution of Blend Strengths in the Subsection of OrchARD's Timbral Augmentation (T.A.) and Timbral Emergence (T.E.) Annotations 19
2.2 Woodwind Instrument Part Distribution for Dominating (Dom.), Embellishing (Emb.), and Emerging (Emer.) Instruments in OrchARD's Blend Annotations . . 20
2.3 Brass Instrument Part Distribution for Dominating (Dom.), Embellishing (Emb.), and Emerging (Emer.) Instruments in OrchARD's Blend Annotations 21
2.4 Percussion Instrument Part Distribution for Dominating (Dom.), Embellishing (Emb.), and Emerging (Emer.) Instruments in OrchARD's Blend Annotations . . 21
2.5 String Instrument Part Distribution for Dominating (Dom.), Embellishing (Emb.), and Emerging (Emer.) Instruments in OrchARD's Blend Annotations 21
2.6 Number of Annotations in which Case-Study Instruments are Featured in OrchARD's Blend Data 24
2.7 Ratio of Case-Study Instruments Dominating versus Embellishing in a Timbral Augmentation Annotation 31

3.1 Association Rule Metrics for the Frequent Patterns in Blend Roles 57
3.2 Association Rule Metrics for the Frequent Patterns in the Violin Case Study . . . 60
3.3 Association Rule Metrics for the Frequent Patterns in the Flute Case Study 62
3.4 Association Rule Metrics for the Frequent Patterns in the Oboe Case Study 63
3.5 Association Rule Metrics for the Frequent Patterns in the French Horn Case Study 65

4.1 Network Statistics 81
4.2 Centrality & Community Detection Results for Dominating Instruments Network . 82
4.3 Centrality & Community Detection Results for the Embellishing Instruments Network 85

4.4 Centrality & Community Detection Results for the Timbral Augmentation Instruments Network . . . 90
4.5 Centrality & Community Detection Results for the Violin Dominating Instruments Network . . . 95
4.6 Centrality & Community Detection Results for the Violin Embellishing Instruments Network . . . 97
4.7 Centrality & Community Detection Results for the Flute Dominating Instruments Network . . . 101
4.8 Centrality & Community Detection Results for the Flute Embellishing Instruments Network . . . 102
4.9 Centrality & Community Detection Results for the Oboe Dominating Instruments Network . . . 105
4.10 Centrality & Community Detection Results for the Oboe Embellishing Instruments Network . . . 107
4.11 Centrality & Community Detection Results for the French Horn Dominating Instruments Network . . . 109
4.12 Centrality & Community Detection Results for the French Horn Embellishing Instruments Network . . . 111

Chapter 1

Introduction

Orchestration as a practice is defined as the selection and combination of different instruments to achieve a specific perceptual result (McAdams, 2019). Often considered to be implicit knowledge gained from practice, orchestration practices and rules have been explained through different composers' treatises. These treatises begin with an examination of the most common instruments found in a Western orchestra, followed by examples and rules for students to learn how to use an orchestra as a musical medium (Piston, 1955). The authors do remain skeptical of a student's ability to effectively learn orchestration through reading and caution the readers accordingly. Berlioz (1948) goes as far as to categorize his treatise as one on instrumentation, which is the selection of instruments based on their acoustic characteristics, rather than on orchestration. He believed that orchestration was a practice that could not be explicitly taught through a set of rules. Rimsky-Korsakov (1912) thought similarly, describing orchestration as an art of creation first and foremost.

There is merit to the argument that orchestration cannot be taught from rules alone. However, this innate orchestration knowledge must have a source. Instincts are born from experience, whether it is accompanied by explicit instruction or not. Music psychologists argue that, by learning these underlying rules through extensive listening and practice, the authors of orchestration treatises are subconsciously forming their rules based on auditory perception principles (Goodchild & McAdams, 2018).

This perspective also allowed for timbre to be introduced as a structuring tool. Timbre is a perceptual attribute that has been historically difficult for musicians and researchers to define. In addition to its perceptual nature, Siedenburg et al. (2019) characterize timbre by its quality and contribution to source identity, its varying granularity of timbral information, and its use in the fusion of auditory events. For these reasons, the theories of timbre and auditory perception piqued the interest of music psychologists, who saw it as their first step to understanding the mystifying instincts at the core of orchestration. McAdams (2019) details timbre's role in orchestration, as well as the increased interest in it in recent years. He connects timbre back to the previously mentioned auditory grouping principles, forming a link between timbre perception and larger perceptual theories of orchestration.

In the process of learning these perceptual theories and applying them to musical examples, it is crucial to understand their underlying patterns. For instance, in orchestral blends, do certain instruments tend to fall into certain functional roles? In these orchestral blends, are there any reoccurring patterns in their instrumentation? When examining a catalogue of musical examples, what can be confirmed or revealed about the orchestration and the instrumentation in these auditory processes? The current study seeks to answer some of these questions regarding the instrumentation of blends. This introductory chapter provides an overview of previous work done in the perceptual theories of orchestration and in instrumentation corpus analysis, in which databases of musical scores are analyzed to search for patterns in their instrumentation. After this background review, the current study and the database of focus will be explained.

1.1 Perceptual Theories of Orchestration

The first perceptual theories to explain the instincts behind the orchestration process were proposed by psychologists who studied auditory cognition and perception. Auditory streams were theorized as a way to explain how listeners connect successive events into musical streams and link them to their respective source. Successive events that had enough continuity in auditory parameters like pitch, loudness, timbre, or texture were considered to be from the same source. If there were to be discontinuity among these parameters, listeners would perceptually organize the sounds as coming from different sources (McAdams & Bregman, 1979). These initial theories, as well as those borrowed from Gestalt psychology, combined to form the auditory scene analysis (ASA) principles. These principles explained the mechanisms surrounding human hearing, with subjects ranging from speech perception to music perception (Bregman, 1990).

Bregman proposed three subcategories within ASA: concurrent groupings, sequential groupings, and segmental groupings. McAdams et al. (2022) defined the three subcategories as follows. Concurrent groupings involve the grouping of co-occurring sounds into distinct musical events. Following this, sequential groupings connect these events into sound streams. Segmental grouping then separates these streams based on their discontinuities to form larger motives and phrases. Music researchers, inspired by Bregman's ASA principles, began to correlate these principles with perceived orchestral effects and musical descriptors. These are outlined in Figure 1.1. In concurrent groupings, listeners can identify certain qualities, such as the timbre, pitch, or spatial position of the musical source. From the connections formed in sequential groupings, different patterns, including rhythmic, melodic, and timbral patterns, begin to emerge. Finally, segmental groupings result in the formation of orchestral contrasts (McAdams et al., 2022).

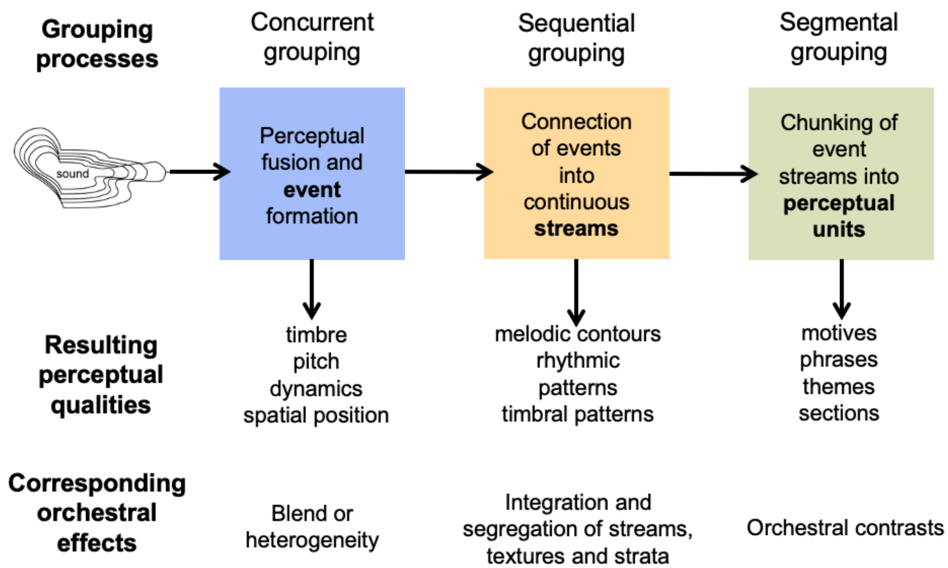

Figure 1.1 Diagram of auditory grouping processes and their related perceptual qualities and orchestral effects. Taken from McAdams et al. (2022)

Based on these descriptions, there is a direct link to be made between auditory grouping processes and orchestration choices. Kendall and Carterette (1991) remark that many musical works of the twentieth century rely on timbral structures as a significant element in their compositions. In fact, as Western music evolved beyond stricter structures and trends, pitch and key became less vital factors for musical structure. Composers began to experiment with pitch-class sets, extended instrument techniques, and tonal ambiguity. This shift is evident in the Straus (2022) score analyses of pieces from the twentieth and twenty-first centuries. In these works, tonal structures are often not the only factor that defines musical form. This makes timbral groupings and organization, which are rooted in listener perception, increasingly attractive for score analysis. Touizrar and McAdams (2019) apply auditory grouping processes in their analysis of Roger Reynolds' *The Angel of Death*. Their work highlights the structure created by auditory perception, as well as timbre's role in creating musical form. Their analysis combines a traditional score analysis with a listening-based analysis. This is of great interest to music theorists, who may add this technique to their tools for analyzing work beyond pitch structures. As such, the role of timbre and auditory grouping processes in orchestration are of great interest to researchers who seek to expand on these music analysis tools.

1.1.1 Instrumental Blend

Figure 1.1 shows that concurrent groupings result in the perception of a blend. A building block for orchestration effects, instrumental blend has long been a subject of research in music cognition. Blend perception is strongly rooted in the Gestalt principle of "common fate". That is, sounds that change in a similar way are more likely to be from the same source and will be perceptually grouped together, resulting in a blend (McAdams, 2019). In auditory perception, this grouping is facilitated by acoustic cues, including onset synchrony, harmonicity, and parallelism in pitch and dynamics (Goodchild & McAdams, 2018). When investigating these cues further, Antoine et al. (2021) remarked that they do not have to be perfect to successfully create a blend. However, deviations in these cues may impact which instrument sounds are grouped into the musical event.

Sandell (1995) coined three terms for blends based on their musical goals and their resulting timbres: timbral augmentation, timbral emergence, and timbral heterogeneity. Timbral augmentation involves two groups of instruments to create a new timbre. The dominating instrument group consists of one or more dominating instruments that are the most prominent in the resulting timbre. The embellishing instrument group encapsulates the rest of the instruments in the blend which often augment the dominating instruments group (McAdams et al., 2022). In other words, dominating instruments are more clearly perceived in a timbral augmentation blend than embellishing instruments. On the other hand, in timbral emergence blends, all instruments involved blend equally to create a new timbre such that the resulting blend does not sound like any of the instruments. The previous two blend categories contrast with timbral heterogeneity, in which all instruments are equally identifiable (Sandell, 1995). McAdams et al. (2022) add that although timbral heterogeneity blends do meet the ASA acoustic cues for grouping, the timbral properties of the instruments allow for each of them to still be heard.

Timbral augmentation and timbral emergence are of great interest as they require deeper knowledge of each instrument's ability to blend. As opposed to timbral heterogeneity, these two blend types are similar to what is taught in many orchestration treatises. Rather than focus on what does not blend, composers and orchestrators focus on what does blend well (Adler, 1982; Berlioz, 1948; Piston, 1955; Rimsky-Korsakov, 1912). The key difference between the two blends is that timbral augmentation possesses a pseudo-hierarchy to the instruments that timbral emergence does not. The dominating instrument group must be carefully constructed such that each instrument maintains its own identity while still incorporating the embellishing instruments into the blend. The embellishing instrument group is also assembled so as to lose their own identity amongst themselves and still augment and highlight the timbre of the dominating group. In this way, timbral augmentation blends appear to be a midpoint between timbral heterogeneity and timbral emergence, which makes them an interesting blend to investigate. As such, they will be the primary focus of this study.

It should be noted that an instrument part appearing as a dominating or embellishing instrument in a timbral augmentation blend does not guarantee that an additional part from that same instrument will have the same role. There are many instances in which two parts from the same instrument are dominating together. However, that is not a guarantee. An example of two instrument parts playing different roles in a timbral augmentation blend can be found in the second movement Beethoven's *Symphony 5.* The first violin is dominating the resulting timbre, whereas the viola, the two bassoon parts, and the second violin are embellishing it. Although both violin parts dominate together elsewhere in this movement, the two parts are different enough in this excerpt for them to not be both dominating the blended timbre. A similar phenomenon can be seen across the woodwind and brass instrument families as well. As such, it is important to keep the instrument parts separate in order to account for examples in which two parts from the same instrument have different roles in a blend.

Nevertheless, understanding blend perception does raise questions on how to effectively create an orchestral blend. What kind of instrumentation is required to achieve a blend of a certain strength? Timbre perception researchers have sought to answer these questions using instrument dyad experiments. In these perceptual experiments, participants are presented with blends consisting of two instruments and asked to rate them by a given scale (Kendall & Carterette, 1991). The controlled environment and the carefully recorded instrument sound samples allow researchers to accurately detect the contributing timbral properties.

The first known experiment of this type was the Kendall and Carterette (1991) perceptual scaling study using dyads of wind instruments. They compared the perceptual properties of instruments playing simultaneously at different harmonic intervals. In their findings, they were able to produce several multidimensional spaces (MDS) using different audio descriptors of the dyads, like nasality, brilliance, and complexity, as the dimensions. This experiment was a first step in analyzing how timbral properties could impact orchestration effects, including blend. The results also outlined how composers could utilize these different sonorities as a musical palette.

Their follow-up study sought to expand their previous findings (Kendall & Carterette, 1993) They conducted instrumental dyad experiments once again, specifically with woodwind instruments, and explored what it meant exactly for instruments to blend. Their findings revealed that instrumental timbre, pitch interval, and time-varying elements, like vibrato, impacted two instruments' ability to blend. Moreover, a stronger blend resulted in the participants' having more difficulty correctly identifying each instrument involved. Although they did not use the notion of timbral augmentation, one might presume that in the case of a timbral augmentation blend, this loss of identity would specifically regard the embellishing instruments. Kendall and Carterette also stressed the importance of scientific blend theories for both acoustic and electroacoustic composers. By understanding how these combinations are made beyond the implicit rules outlined in orchestration treatises, artists can have more control over their instrumentation choices.

Sandell (1995) used dyads of computer-generated instrument tones to achieve different blend strengths. The sounds were created by Grey (1977) from instrument tone recordings that were analyzed and resynthesized in a simplified version. Sandell found that spectral centroid, which is an audio descriptor associated with an instrument's brightness, was fundamental to an instrument's ability to blend. Specifically, higher spectral centroids have lower perceived blend scores, indicating that "darker" instruments are easier to blend than "brighter" instruments. As Sandell explains, this phenomenon is in line with what has been written by composers in orchestration treatises (Adler, 1982; Piston, 1955; Rimsky-Korsakov, 1912). However, he also remarks that instrument combinations with vastly different spectral centroids had the weakest perceived blend. From a composer's perspective, these findings can further explain certain instinctive orchestration choices.

Another study involved dyads of different sustained and impulsive instrumental sounds with the aim of comparing the two sound types (Tardieu & McAdams, 2012). The perceived blend ratings indicated that both sound types blend more effectively if they have low spectral centroid values. In the case of impulsive sounds, the ones with a slower attack time had higher blend scores from participants. Ultimately, their attack time had a greater influence since impulsive sounds have less time to resonate than their sustained counterparts. Considering that onset synchrony is crucial to concurrent grouping (Antoine et al., 2021), Tardieu and McAdams conclude by mentioning that impulsive sounds have greater influence on blend perception, whereas sustained sounds have greater control on the overall timbre.

In general, scientific research into blend is of great interest for composers and music psychologists alike, with the shared aim of understanding orchestration on a deeper level. Instrument dyad studies have given greater insight into the timbral descriptors that affect instrumentation choices. They may also explain why certain instruments tend to fall into certain functional roles in an orchestral work.

1.2 Previous Work: Instrumentation Corpus Studies

When asking these larger questions about the orchestration of blends, it is important to consider many different examples to provide accurate patterns. Corpus studies are a powerful research tool used in many disciplines to analyze large databases. In music cognition, corpus research can be used to gather different listener experiences and find patterns, which can eventually be used to create predictions and models. It is also used in music research to detect instrumentation patterns of certain composers or eras. Temperley and VanHandel (2013) highlighted major contributions to music corpus research, ranging from musicological applications to cognitive approaches for composition. However, it is important to note that all corpora have a certain degree of limited representation by nature. As such, many corpus studies focus on specific questions, either fixating on instrumentation evolution or on instrument roles in a specific epoch (Le et al., 2022).

Johnson (2011) analyzed a corpus consisting of nineteenth-century orchestral symphonies with the aim of finding a middle ground between the instrument family and the timbral contrasts to explain certain instrumentation choices. In Johnson's case, corpus analysis uncovered a different perspective to instruments' purpose in orchestration. Rather than fixate solely on timbral properties or orchestration restrictions, a new model was created from an instrument's frequency of occurrence. To create his chord database, Johnson selected 50 works from a total of 230 from which to build a model, sampling a single sonority from each selected piece. He then rated each possible instrument pairing on a scale from 0.0 to 1.0, with a higher score indicating a higher chance that both instruments would play at the same time. These distances were mapped onto an MDS model. The use of MDS to characterize timbre spaces was first introduced by Plomp (1970). The meaning of each dimension, while not initially defined, is extrapolated through the researcher's interpretation. However, Johnson successfully identified only one dimension: the instrument's dynamic potential. Johnson thus had to conduct an additional hierarchical clustering analysis to further explore the differences between the instrument clusters. The distances between the data points created three clusters, forming the Standard, Power, and Colour (SPC) model. Each cluster was characterized by the instrument's different roles in orchestration. The Standard instruments performed the most often and were used in both melodic and accompaniment roles. The Power instruments were dynamically intense and covered the largest pitch range. The Colour instruments performed quietly and less often than other instruments.

Le et al. (2022)'s instrumentation study used an orchestral music corpus to explore instrument roles across different textural layers. In this case, they chose to focus on the streams formed in sequential grouping. The corpus consisted of annotations from the first movements of Classical era orchestral symphonies. After defining four different layer roles (melodic, rhythmic, harmonic, mixed), the instrument distribution across each layer was calculated and analyzed to determine specific roles and instrument combinations between the layers. It was discovered that instruments of different families rarely appeared in the same layer at the same time. In fact, alternating between different families was found to be a common strategy to create call-and-response effects. Concerning roles, strings were used most often for melody, whereas brass and woodwinds were mostly used in harmonic roles. The only percussion instrument featured, the timpani, was primarily used in rhythmic and mixed layers. This corpus study is particularly distinct from other presented in this section as it applies the strategy of instrumentation research to sequential groupings. Johnson (2011), on the other hand, made it clear that he wanted to develop a model distinct from those concepts.

Chon (2022)'s literature review of instrumentation corpora studies includes analyses that focused on other musical descriptors, like pitch, dynamics, and tempo, rather than timbre and blend. For example, one was able to examine the presence of different instruments over the years, as certain instruments were invented and others fell out of fashion (Chon et al., 2018). Chon concludes the literature review by highlighting that these evolutionary corpus studies, while confirming obvious intuitions in instrumentation, can reveal overlooked patterns.

The common factor between all these instrumentation corpus studies is the search for instrument roles in orchestration. Each worked within the data limitations of its corpus, such as Le et al. (2022) restricting themselves to symphonies of Haydn, Mozart, and Beethoven. Some, such as Johnson (2011) even constructed their corpus with the expressed intent of avoiding overrepresentation of one composer or one piece. Overall, they all confirmed the value of instrumentation corpus research, both to support previous findings and to point out potentially missed patterns.

1.3 The Current Study

1.3.1 Objectives

Auditory grouping processes and instrumental blend studies have been crucial to understanding the underlying perceptual theories of orchestration. Through examining the contributing factors to blend in closed environments, we can discover to what degree two sounds blend effectively together prior to manipulation from the composer, the performer, the conductor, or the recording engineer (Kendall & Carterette, 1993). As discussed earlier in this chapter, the timbral descriptors affecting blend have been explored at length through perception experiments. However, there is a benefit in using auditory grouping processes as a music analysis technique, especially for more recent musical works (Kendall & Carterette, 1991; Touizrar & McAdams, 2019). After all, orchestration effects and timbral contrasts provide a structure just as much as chordal structures do (McAdams, 2019). It is therefore of great interest to study how the current timbral blend theories compare to instrumental music examples, and if there are any reoccurring patterns in the blends' instrumentation.

The primary question guiding the current study involves the patterns within instrumental blends. Do some instruments fall into certain functional roles in orchestral blends? Some orchestration treatises, such as the Rimsky-Korsakov (1912) treatise, dedicated sections to blends and provided examples of instrument combinations that resulted in what he considered to be successful blends. Others have offered examples in which certain instruments compliment the tone colour of other instruments (Adler, 1982; Piston, 1955). Without directly using the terms timbral augmentation or timbral emergence, they indirectly address methods to achieve these effects. However, these treatises cannot be directly applied to timbral augmentation blends. It can be hypothesized that instruments that are often used in the melody line, such as violins, may also be used as dominating instruments. Inversely, instruments that often assume harmonic or "colour" roles, as outlined in treatises, may be used more often as embellishing instruments. The current study aims to answer this question and to explore these theories.

The relationships between these instruments are also considered. Do certain instruments dominate together or embellish together? Moreover, when a given instrument dominates a timbral augmentation, which instruments embellish it? These questions are, once again, somewhat addressed in orchestration treatises when composers discuss certain instrument pairings. For instance, Adler (1982) mentions the tonal balance created when a clarinet, flute, and oboe play at the same time. Does a similar relationship exist in the different functional roles of a timbral augmentation blend?

The current study uses a corpus analysis to explore the instrumentation patterns and relationships in timbral augmentation blends. This type of analysis was selected in the interest of exploring blends across several different pieces and composers. Of the previously explored corpus studies, many focused on the individual roles of the instruments, with some exceptions such as Le et al. (2022)'s textural layer analysis. However, blends are influenced by the relationships between instruments. As such, the analysis techniques will be carefully selected to take the instrument interactions into account. Some methods will be used to discover which instruments are more frequently used in certain functional roles (see Chapters 2 and 3). Others will explore some common combinations and the relationships between these instruments (see Chapters 3 and 4). All analysis techniques used will consider timbral augmentation blends of different perceived strengths, as well as timbral emergence blends so as to differentiate embellishing instruments from those used in timbral emergences. These findings will be compared to orchestration treatises to bridge the gap between auditory perception research and composers' intuitive theories. The orchestration treatises used were Adler (1982), Berlioz (1948), Piston (1955), and Read (1979). The four texts had been previously analyzed for mention of auditory grouping principles by another project in the Music Perception and Cognition Lab, and the sections relating to blend were noted. In addition, they focus extensively on works from the Classical and Romantic eras, which are the most prevalent eras featured in the data analyzed for the current study. Although the other three authors wrote their treatises long after these musical epochs, Berlioz, a Romantic composer, wrote his in his lifetime based on his own work. For these reasons, these four texts were the ideal resources for the current study.

1.3.2 The Orchestration Analysis & Research Database (OrchARD)

The Orchestration Analysis & Research Database (OrchARD)[1] is a database consisting of auditory grouping process annotations found in orchestral music. OrchARD was created by the ACTOR Project's Orchestration and Perception Project. The mission of this group is to create a psychological theory of orchestration based on different perceptual principles, such as ASA or auditory grouping processes, rooted in timbre.[2] OrchARD was created with the aim of building this orchestration theory by amassing score analysis examples (Russell et al., 2016). This corpus and its annotations were the subject of the current study.

The annotations were created by a team of experts in the field of auditory grouping processes. Each movement or piece was selected from one of 64 movements across 24 different pieces and was then assigned to two experts. For every excerpt, the two experts were asked to identify the auditory grouping process, its resulting perceptual effect, and the strength of the effect. Although the experts had access to the scores for their annotations, their decisions were made from listening to a specific recording of the piece. Their initial annotation process was done individually, and they were instructed to maintain a journal detailing any information related to their decision-making process (Russell et al., 2016). Following the individual annotations, the two assigned experts met to discuss their analyses. If they agreed on their annotation, it was accepted. If they disagreed, they discussed their perspectives with the goal of reaching an agreement. If they were still unable to arrive at a consensus, they turned to the larger team of music theorists, composers, and psychologists to give the final verdict.

The annotations were initially kept in spreadsheets. However, this data storage method was cumbersome and difficult to navigate as more data was added. As such, OrchARD was created to provide researchers with a more flexible and accessible tool with which the database contents could be queried (Russell et al., 2016). Figure 1.2 shows a sample annotation from Schubert's Symphony 9. As can be seen in this example, the appropriate page of the score is included with the annotated section clearly marked. Each annotation also includes the excerpt's auditory grouping process, the resulting perceptual effect, and the effect's strength, as annotated by the experts. The annotations also include the metadata related to the excerpt's movement, composer, and recording.

The effect's duration is also noted either as punctuated or sustained. Punctuated means that it is short, such as one note played by the whole orchestra *in tutti*. Sustained, on the other hand, implies that the effect is continuous and can either be stable, meaning that the instruments involved remain the same throughout the excerpt, or transforming, meaning that they change over

---

[1] https://orchard.actor-project.org/

[2] https://sites.music.mcgill.ca/orchestration/

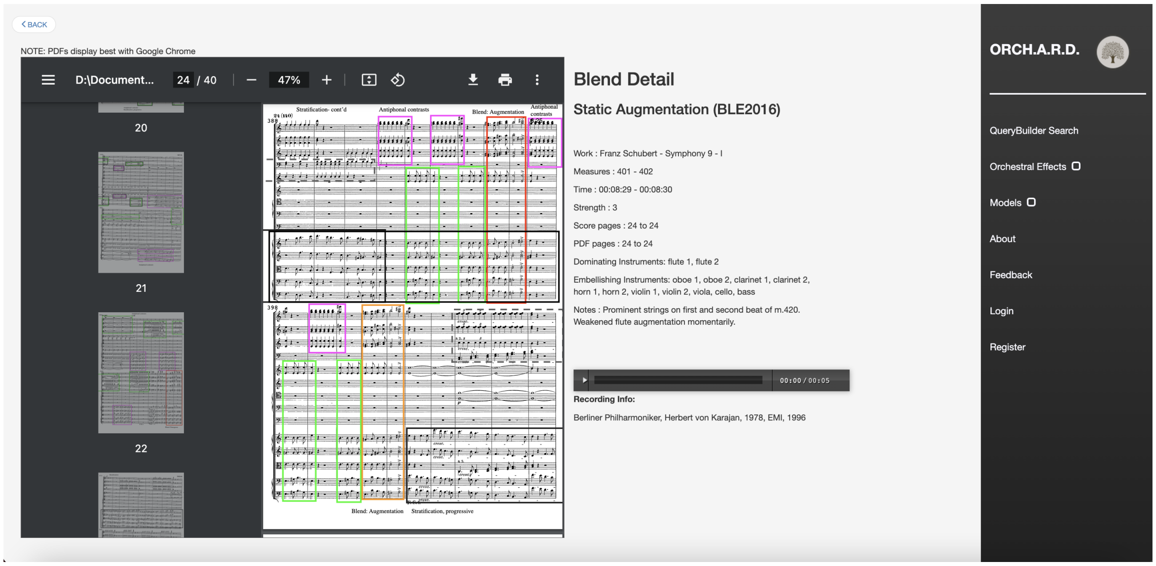

Figure 1.2 Example of OrchARD annotation as seen on the current OrchARD user interface. The different colours on the score's page reflect different types of orchestral effects found in the passage. The timbral augmentation shown here is represented by the orange box on the bottom half of this page.

time (McAdams et al., 2022). The analysis techniques used in the current study do not directly take into account the effect's duration or its label of stable or continuous. However, they will be considered when interpreting some of the results.

Rather than examine all auditory grouping processes, the current study is focused on timbral augmentation blends and timbral emergence blends. It should be noted, however, that OrchARD contains more annotations for timbral augmentation blends than for timbral emergence blends. Table 1.1 shows the number of timbral augmentation and timbral emergence blend instances for each composer featured in OrchARD. These numbers will be further analyzed in Chapter 2, and the resulting limitations of this discrepancy will be discussed in Chapter 5.

Table 1.1: Number of blend entries in OrchARD by composer

Composer	Number of Timbral Augmentation Blend Entries	Number of Timbral Emergence Blend Entries
Wolfgang Amadeus Mozart	11	7
Joseph Haydn	65	10
Ludwig van Beethoven	34	25
Franz Schubert	156	72
Felix Mendelssohn	66	10
Modest Mussorgsky	244	79
Giuseppe Verdi	25	4
Georges Bizet	9	3
Alexander Borodin	17	0
Bedřich Smeltana	26	16
Johannes Brahms	119	10
Gustav Mahler	68	21
Jean Sibelius	34	6
Vincent d'Indy	8	0
Claude Debussy	92	38
Ralph Vaughan Williams	100	12
Frank Ticheli	94	58

Table 1.2 details the pieces and recordings found in the annotations of the subset in OrchARD. As of September 2022, when the data was retrieved, some of the recording entries in OrchARD were incomplete. This was largely the case for the recording years, some of which were not specified in the database. For this reason, the *Recording Details* cell contains the orchestra, the conductor, and the year of the recording if available.

Table 1.2: Pieces and movements featured in OrchARD's timbral augmentation and timbral emergence blend data

Composer	Title of Piece	Movement(s)	Year of Composition	Recording Details (Orchestra, Conductor, Year if Available)
Wolfgang Amadeus Mozart	Don Giovanni	Overture	1787	Berliner Philharmoniker, Daniel Barenboim
Joseph Haydn	Symphony 100	I-IV	1793	The Academy of Ancient Music, Christopher Hogwood
Ludwig van Beethoven	Symphony 5	I-II	1808	Tonhalle Orchester, David Zinman
	Symphony 7	I	1812	Berliner Philharmoniker, Herbert von Karajan
Franz Schubert	Symphony 8	I-II	1822	Lucerne Festival Orchestra, Claudio Abbado
	Symphony 9	I-IV	1825-1828	Berliner Philharmoniker, Herbert von Karajan, 1978
Felix Mendelssohn	Symphony 3	I-IV	1829-1842	London Symphony Orchestra, Peter Maag, 1957
Modest Mussorgsky	Symphonie Fantastique	I-IV	1830	Orchestre Symphonique de Montréal, Charles Dutoit, 1984
	Pictures at an Exhibition	1-15	1874	Berliner Philharmoniker, Simon Rattle, 2007
Giuseppe Verdi	Rigoletto	Prelude	1851	Italiana Opera Orchestra, Georg Solti
	La Traviata	Overture	1853	Philharmonia Slavonica, Henry Adolph
	Aida	Overture	1871	Wiener Philharmoniker, Herbert von Karajan
Georges Bizet	Carmen	Overture	1875	Orchestre Symphonique de Strasbourg, Alain Lombard
Alexander Borodin	In the Steppes of Central Asia	N/A	1880	USSR State Symphony Orchestra, Evgeny Svetlanov
Bedřich Smeltana	Ma Vlast	II	1880	Berliner Philharmoniker, Herbert von Karajan
	The Bartered Bride	Overture	1886	Vienna Radio Symphony Orchestra, Alfred Scholz (Most Famous Opera Overtures)
Johannes Brahms	Symphony 4	I-IV	1884-1885	Berliner Philharmoniker, Herbert von Karajan, 1978
Gustav Mahler	Symphony 1	I-IV	1884-1888	Bavarian Radio Orchestra, Rafael Kubelik, 1969
Jean Sibelius	Symphony 2	II	1902	London Symphony Orchestra, Sir Colin Davis
Vincent d'Indy	Choral Varié	N/A	1903	Iceland Symphonic Orchestra, Rumon Gamba
Claude Debussy	La Mer	I-III	1903-1905	Berliner Philharmoniker, Simon Rattle, 2001
Ralph Vaughan Williams	The Lark Ascending	N/A	1920	London Philharmonic Orchestra, Andrew Litton
	Symphony 8	I,II,IV	1943	London Philharmonic Orchestra, Adrian Boult
Frank Ticheli	Symphony 2	I-III	2003	Dallas Wind Symphony, Jerry Junkin, 2011

It should also be noted that OrchARD's data possesses a limited representation of musical styles and epochs. The musical works are largely taken from the Classical and Romantic eras, with Ticheli's *Symphony 2* being the only piece from the twenty-first century. *Symphony 2* is also the only piece composed for concert band in the entire dataset. As such, its instrumentation was vastly different from the other pieces in the database, which were written for orchestra. The annotations associated with Ticheli's piece were therefore removed pre-emptively from the data. The lack of representation and its implications will be discussed in more detail in Chapter 5 (see section 5.2).

At the time of this study, OrchARD is semi-restricted to public access. An older version of the database is currently available upon request from the Orchestration and Perception Project. As a result, its previous research applications are minimal. OrchARD's annotations were used as a ground truth when building an orchestral effect prediction model (Antoine et al., 2021). Using symbolic data extracted from scores, the model was trained to label orchestral effects. The prediction was based on calculations of the onset synchrony, the harmonicity, and the parallelism in pitch and dynamics of a certain excerpt. The model's accuracy was then tested using the excerpts found in OrchARD. The dataset's annotations were then used as ground truth for the evaluation stage. This model was ultimately able to achieve 81% accuracy, with some significant errors occurring in the nearly 20% of mislabellings. Beyond its involvement in the model's testing phase, OrchARD's data has not been the focal point of any other research prior to the current study.

1.3.3 Overview

The current study applied three analytical approaches to OrchARD's data. Each technique answered a different question regarding the orchestration of timbral augmentation blends. The first analysis technique explored the relative frequencies of occurrence for each dominating and embellishing instrument in the data. Though simple in its approach, the aim of this method was to give an initial glance into the combinatorial patterns found in the blend data, which was then explored in greater depth in subsequent analyses. In addition, the relative frequencies of occurrence in the timbral emergence data were considered. This allowed for a comparison between the two blend types and determined if there are any differences between the different blend roles (i.e., dominating instruments, embellishing instruments, emerging instruments). The analysis methodology and results are outlined in Chapter 2.

The second analysis technique applied frequent pattern mining to the dataset so as to explore larger frequently occurring groupings, as well as the different frequent patterns at different strengths of blend. Developed initially for market analysis, the objective of frequent pattern mining is to discover correlations and build associations among the items found in the dataset (Han et al., 2007). In the current study, the items investigated were instrument sections found within OrchARD's annotations. Following the previous inspection of simple pairings and instrument distribution, frequent pattern mining was used to expand on those patterns and to provide more insight into them through associations. Background information on frequent pattern mining, as well as the experiment's methodology and results, will be discussed in Chapter 3.

The third and final analysis technique was a network analysis of the timbral augmentation data. An emerging field in the 21st century, network science examines complex data and their relationships in a computational and visual medium (Barabási, 2016). In contrast to the previous two experiments, this procedure inspected the relationships and functional roles performed by instrument sections when they were either the dominating instrument or the embellishing instrument. The goal was to explore the relationships between the instruments in the timbral augmentation blend role and explore potential subgroupings within each blend role. Chapter 4 provides a brief summary of network science concepts used in the experiment before explaining the methodology and results of the different network analysis algorithms applied.

Chapter 5 summarizes all three analysis techniques, relating their results to what has been written in orchestration treatises. This returns to the overall objective of the current study. The limitations of this analysis are addressed, and future research directions are proposed. Little research has been conducted using OrchARD's data as of yet, so there is still much to be explored.

Chapter 2

Experiment 1: The Relative Frequencies of Occurrence of Different Types of Blend

The first step when conducting a corpus analysis with a relatively unexplored database is to familiarize oneself with its general characteristics. Although this initial examination may not directly reveal patterns in larger datasets, it could inspire preliminary theories regarding the data's behaviour, which can then be further tested in follow-up analyses. Although OrchARD was briefly introduced in Chapter 1, its blend data were not explored in depth.

As we begin to ask questions about the patterns in instrumentation in timbral augmentation blends, it is important to establish a clear distinction with timbral emergence blends. Beyond their different functions in orchestration, as described by Sandell (1995), are there ways to distinguish them based on their instrumentation? Furthermore, the definitions of embellishing instruments in timbral augmentations and instruments in timbral emergences indicate that they share a common function: to blend effectively (McAdams et al., 2022). Do they have similar instrumentation as a result of this shared goal?

This chapter encompasses the first step of the current study's corpus analysis. The relative frequencies of occurrence for each instrument part featured in the timbral augmentation and timbral emergence blend data were calculated to detect early patterns in each blend role. Timbral emergence data were also considered for this analysis so as to create a contrast between this blend type and timbral augmentation blends. This data subset's preliminary statistics will be reported first to give an initial overview. Following that first glance, we will discuss the methodology and results of this analysis and compare the findings to the instrumentation suggestions offered by orchestration treatises.

2.1 Method

2.1.1 Basic statistics of dataset

At the start of this analysis, OrchARD's data subset was examined to obtain preliminary information on timbral augmentation and timbral emergence blends. As mentioned in Chapter 1, there are a total of 64 movements taken from 24 different pieces. There are also a total of 90 different instrument parts that can be possibly featured in an annotation. These instrument parts fall into one of four different instrument families: woodwind, brass, percussion, and strings. OrchARD has an additional family category, "other", which encompasses instruments that seemingly do not fit as neatly into the four families, such as keyboard instruments (celesta and organ) and plucked strings (harp). For this analysis, the "other" family was removed, and its instruments were added to whichever family suited it the best. The celesta was added to the percussion family, the organ was added to the woodwind family, and both harp parts were added to the string family.

Although the primary focus of the current study is timbral augmentation blends, timbral emergence blends were also considered in this first analysis. This was done to form a comparison between the instrumentation distribution of the two blends. In theory, the embellishing instruments in timbral augmentations should have many similarities to the instruments found in timbral emergences. Instruments in both roles must blend effectively and not be extremely noticeable in the resulting timbre. However, timbral emergences require that all instruments involved blend together cohesively, whereas embellishing instruments are meant to blend together while enhancing the timbre of the dominating instruments. This slight difference may have an impact on the instrumentation in both roles. To answer this with more certainty, timbral emergence blends had to be considered.

The initial counts presented in this chapter focus exclusively on the data subsets of timbral augmentation blends and timbral emergence blends. Table 2.1 shows the blend strength distribution for both types. Figures 2.1 and 2.2 visually represent these distributions. The initial shape from the bar graphs suggests that the blend strength distributions for timbral augmentation and timbral emergence blends follow a normal distribution. After conducting a Shapiro-Wilk test, both sets of data failed to reject the null hypothesis that they follow a normal distribution at the 5% significance level, thus supporting this observation. An examination of these values reveals that the blend strengths follow a normal distribution. This is most apparent with the timbral augmentation blends. Although the timbral emergence has a slight normal distribution, its peak is not at the center value (i.e., 3), but rather at 4.

Table 2.1: Distribution of Blend Strengths in the Subsection of OrchARD's Timbral Augmentation (T.A.) and Timbral Emergence (T.E.) Annotations

Blend Strength	# of Instances (T.A.)	# of Instances (T.E.)
1	44	13
2	236	38
3	422	74
4	298	122
5	74	56

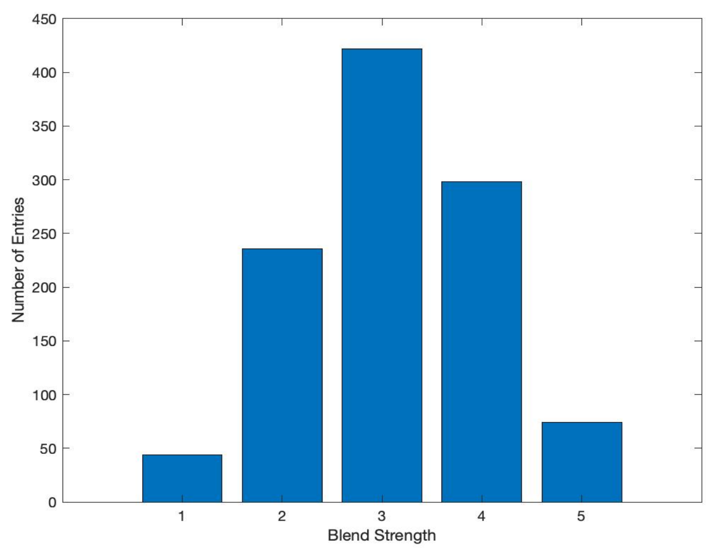

Figure 2.1 Blend Strength Distribution for Timbral Augmentation Blends

Tables 2.2, 2.3, 2.4, and 2.5 provide the instrument part counts for the woodwind, brass, percussion, and string families respectively. These four tables show the difference between the distributions across all blends for each blend role. Instrument parts which were not represented in any blend role were omitted.

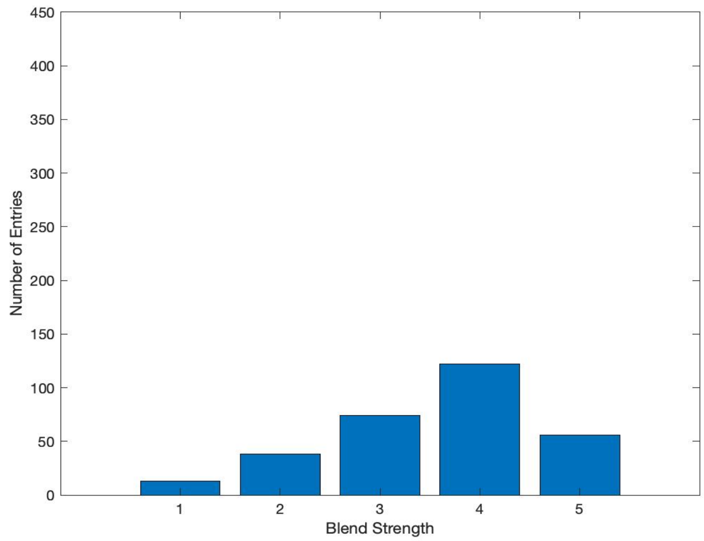

Figure 2.2 Blend Strength Distribution for Timbral Emergence Blends

Table 2.2: Woodwind Instrument Part Distribution for Dominating (Dom.), Embellishing (Emb.), and Emerging (Emer.) Instruments in OrchARD's Blend Annotations

Instrument Part	# of Instances (Dom.)	# of Instances (Emb.)	# of Instances (Emer.)
Piccolo	26	82	47
Flute	282	873	407
Oboe	153	999	454
English Horn	7	56	47
Clarinet	84	1157	454
Piccolo Clarinet	3	13	16
Bass Clarinet	1	52	34
Bassoon	59	1120	455
Contrabassoon	3	101	37
Alto Saxophone	3	3	8

Table 2.3: Brass Instrument Part Distribution for Dominating (Dom.), Embellishing (Emb.), and Emerging (Emer.) Instruments in OrchARD's Blend Annotations

Instrument Part	# of Instances (Dom.)	# of Instances (Emb.)	# of Instances (Emer.)
Horn	429	938	516
Bass Horn	0	11	0
Trumpet	213	267	254
Cornet	17	8	16
Trombone	120	176	172
Alto Trombone	17	32	35
Bass Trombone	18	39	32
Euphonium	1	0	0
Tuba	10	41	40

Table 2.4: Percussion Instrument Part Distribution for Dominating (Dom.), Embellishing (Emb.), and Emerging (Emer.) Instruments in OrchARD's Blend Annotations

Instrument Part	# of Instances (Dom.)	# of Instances (Emb.)	# of Instances (Emer.)
Timpani	8	179	120
Triangle	0	17	16
Tam-Tam	0	2	16
Gong	0	1	2
Cymbals	0	20	41
Rattle	0	0	5
Whip	0	0	5
Ratchet	0	0	1
Snare Drum	0	5	11
Bass Drum	0	24	30
Tambourine	0	0	6
Glockenspiel	0	7	6
Tubular Bells	0	2	15
Vibraphone	1	0	0
Xylophone	0	6	8
Celesta	1	2	6

Table 2.5: String Instrument Part Distribution for Dominating (Dom.), Embellishing (Emb.), and Emerging (Emer.) Instruments in OrchARD's Blend Annotations

Instrument Part	# of Instances (Dom.)	# of Instances (Emb.)	# of Instances (Emer.)
Harp	9	83	63
Violin	575	322	297
Viola	185	241	152
Cello	192	253	144
Bass	124	242	142

After removing the outlier piece mentioned in Chapter 1, OrchARD contains a total of 1,074 timbral augmentation annotations and 313 timbral emergence annotations. A timbral augmentation blend involves on average 2.37 dominating instruments and 6.87 embellishing instruments, which is a total of approximately 9.23 instruments on average. The resulting blend has an average strength of 3.11. In comparison, timbral emergence blends involve an average of 13.31 instruments, resulting in an average blend strength of 3.44. These initial statistics indicate the key difference between timbral augmentations and timbral emergences. There are more instruments involved in creating a timbral emergence blend than a timbral augmentation blend, and the resulting blend is stronger.

Based on the results from the previous Shapiro-Wilk test for normal distribution within the blend strength distributions, another normality test was conducted using this data. An independent-groups t-test was conducted with the null hypothesis that the average number of instrument parts and the average blend strengths for the two blend types are independent random variables taken from two normal distributions. The resulting p-value was at 0.06, which is marginally significant enough to not reject the null hypothesis at 5% significance. This, however, is in stark contrast to the previous Shapiro-Wilk results, which produced p-values of 0.66 and 0.91 for timbral augmentation and timbral emergence blends, respectively. In short, the average number of instrument parts for each blend and the average blend strength for each blend weakly follow a normal distribution.

2.1.2 Procedure

The objective of the following analysis was to become familiar with the content of OrchARD's timbral augmentation data. To do so, the relative frequencies of occurrence for each instrument family, each instrument, and each instrument pairing with certain case-study instruments were calculated. The term "instrument pairings" refers to two instruments that are featured in the same blend annotation. This could include, for instance, two instruments in the same timbral emergence blend or two instruments in the same timbral augmentation blend, either in the same blend role or in different ones. The subsequent analyses in the upcoming chapters will then be able to build on and clarify certain points missed by these simpler calculations.

Relative frequencies of occurrence were chosen as the first analysis technique in order to place each instrument's use in the context of its blend role. The absolute counts in Tables 2.2–2.5 do provide some insight. However, these counts do not portray an accurate assessment of how each instrument is used in each blend role. As mentioned previously, each blend role has a different average number of instruments identified in each annotation and a different total number of instrument instances. Relative frequency scales these absolute values to a range from 0 to 1, based on how frequently they occur in the context of how many instances of each instrument there are in a given blend role. Such a technique makes it easier to compare how often each instrument occurs in the three blend roles.

To calculate the relative frequencies of occurrence for each blend role, the number of instances of each instrument across all its parts was divided by the total number of instrument part instances in that blend role. In total, there are 2,542 dominating instrument part instances, 7,374 embellishing instrument part instances, and 4,110 emerging instrument part instances. The relative frequencies of each instrument were visualized using bar graphs. In addition, a pie chart was created for each of the blend roles to show the distribution of each instrument family in each role.[1]

Case studies

Certain instruments were selected for case studies in the following analysis to gain a deeper understanding of specific instruments. Orchestration treatises dedicate sections to individual instruments rather than attempting to discuss blends on a large scale. So, case studies were used to achieve a similar level of familiarity and compare instruments from different families. The case-study instruments were chosen according to the preliminary OrchARD statistics outlined in the previous section. The sole criterion in the selection process was the instrument's value in the dominating instruments column in Tables 2.2–2.5. As such, violin, flute, oboe, and French horn (which will be used interchangeably with horn) were chosen as our case-study instruments.

To calculate the relative frequencies of occurrence for the case-study instruments, a slightly different technique was needed to get the total number of annotations in the denominator. Annotations were counted on the basis of whether or not they included one of the case study instruments' parts. If an annotation possessed more than one part of a given instrument, it was not counted multiple times for each part present. For example, if oboe 1 and oboe 2 were listed as embellishing instruments in an annotation, it would be considered as one instance of an oboe as an embellishing instrument. These counts are outlined in Table 2.6. For timbral augmentation blends, only the annotations in which a given instrument is a dominating instrument were counted. This technique also extended to the instrument pairings, meaning that the resulting bar graphs will not be divided into an instrument's separate parts as they were in the instrument role counts.

[1] The calculations and graph generation for this chapter were conducted using MATLAB code. The file is available upon request from stephen.mcadams@mcgill.ca.

Table 2.6: Number of Annotations in which Case-Study Instruments are Featured in OrchARD's Blend Data

Instrument	# of Annotations (T.A.)	# of Annotations (T.E.)
Violin	361	155
Flute	173	203
Oboe	94	227
French Horn	157	169

2.2 Results

2.2.1 Analysis of blend roles

Dominating instruments

Figure 2.3 shows the instrument family distribution for dominating instruments. Each family is colour-coded, with dark yellow representing woodwinds, green representing brass, orange representing percussion, and pink representing strings. This colour scheme remains consistent across the graphs present in this chapter. Figures 2.4–2.7 show this distribution across all four families, now by instrument.

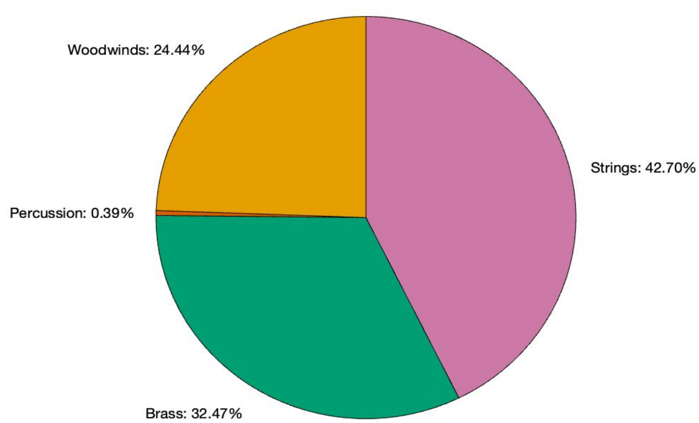

Figure 2.3 Distribution of Instrument Families in Dominating Instruments

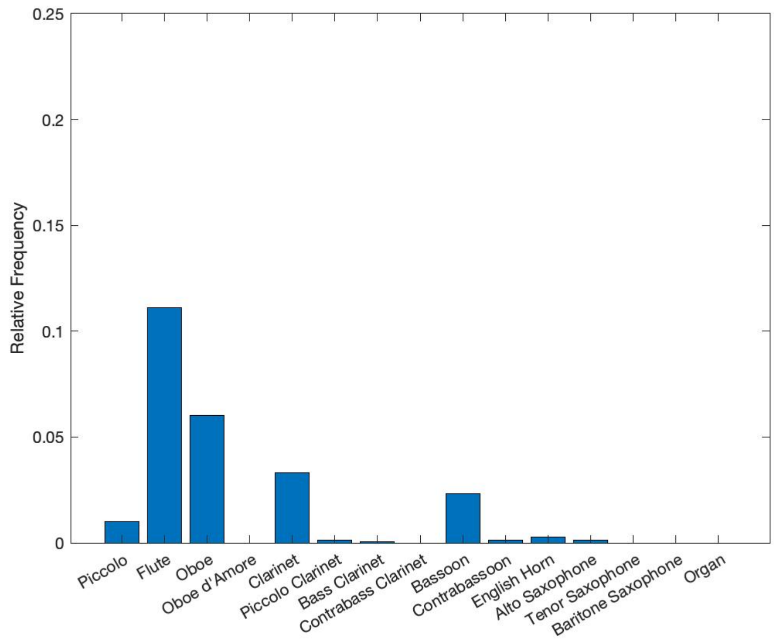

Figure 2.4 Distribution of Woodwind Instruments in Dominating Instruments

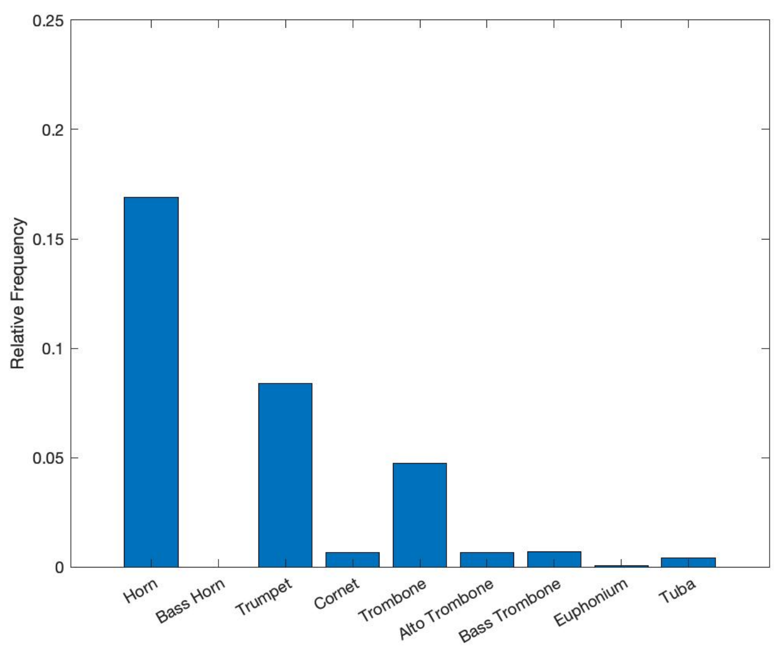

Figure 2.5 Distribution of Brass Instruments in Dominating Instruments

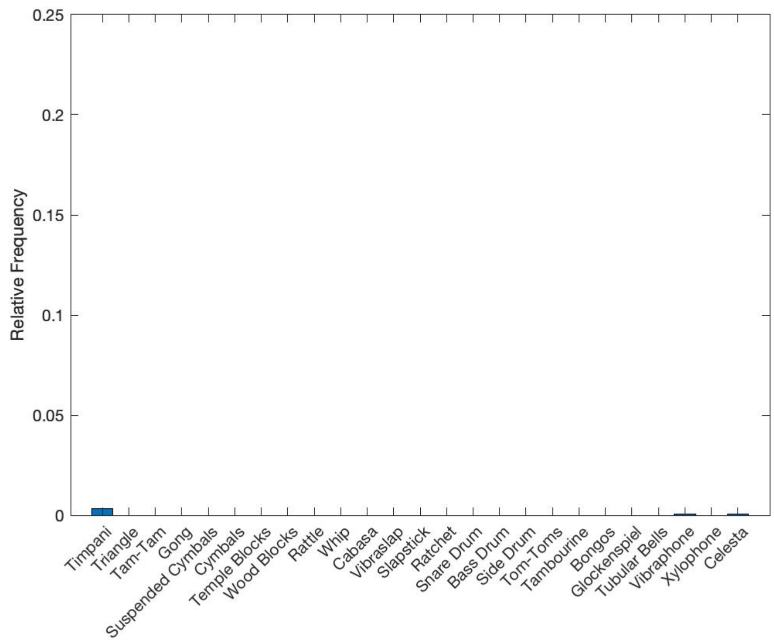

Figure 2.6 Distribution of Percussion Instruments in Dominating Instruments

Figure 2.1 dissects the frequency with which certain instrument families are used. Woodwinds only account for 24.44% of dominating instrument instances. Figure 2.2 shows that the dominant instruments that are most frequently used in the woodwind family are flutes and oboes. Bassoons and clarinets have some representation; however, their involvement is comparably lower than the aforementioned two instruments. The brass family's segment on the pie chart (Figure 2.1) is also comparably large, as the family is found in 32.47% of dominating instrument instances. As shown in Figure 2.3, horns and trumpets are the most frequently used dominating instruments in their family, with horns being used more often than trumpets. The percussion family is rarely present as dominating instruments, as shown in Figure 2.4. The most frequently used instrument is the timpani, comprising 0.3% of dominating instrument instances. Compared to the other three families, strings are the most frequently used dominating instruments, accounting for 42.70% of dominating instrument instances. According to Figure 2.5, violins possess the highest relative frequency of occurrence of all dominating instruments in this dataset.

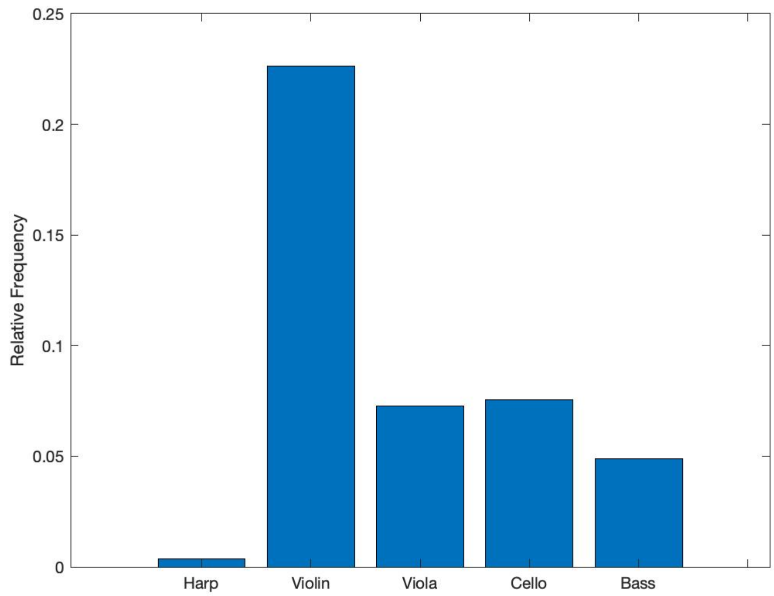

Figure 2.7 Distribution of String Instruments in Dominating Instruments

Embellishing instruments

Figure 2.8 shows the instrument family distribution for the embellishing instrument blend role. Figures 2.9–2.12 then show the instrument distribution across the four instrument families. These five graphs were constructed in a similar manner to the previous blend role's figures.

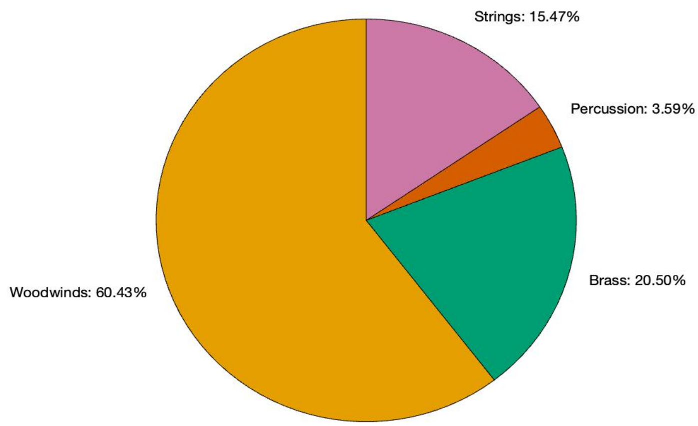

Figure 2.8 Distribution of Instrument Families in Embellishing Instruments

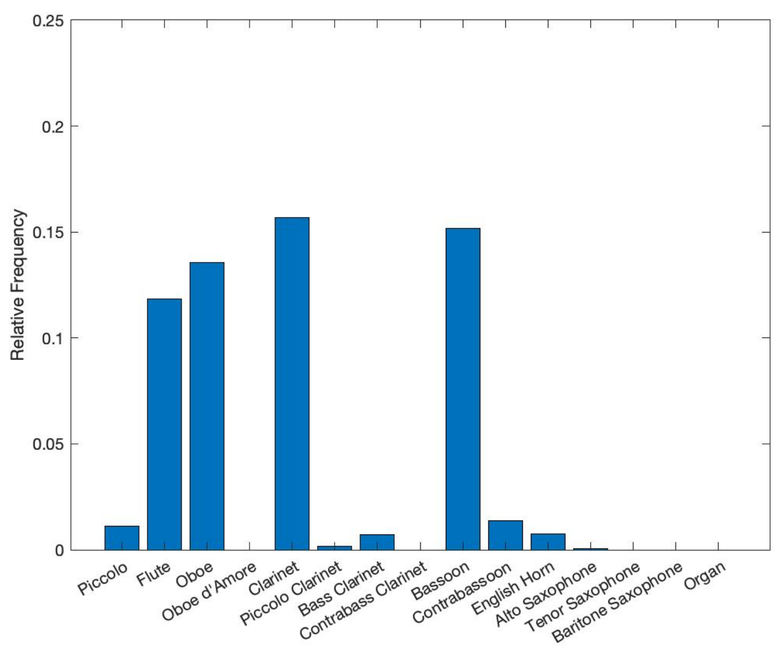

Figure 2.9 Distribution of Woodwind Instruments in Embellishing Instruments

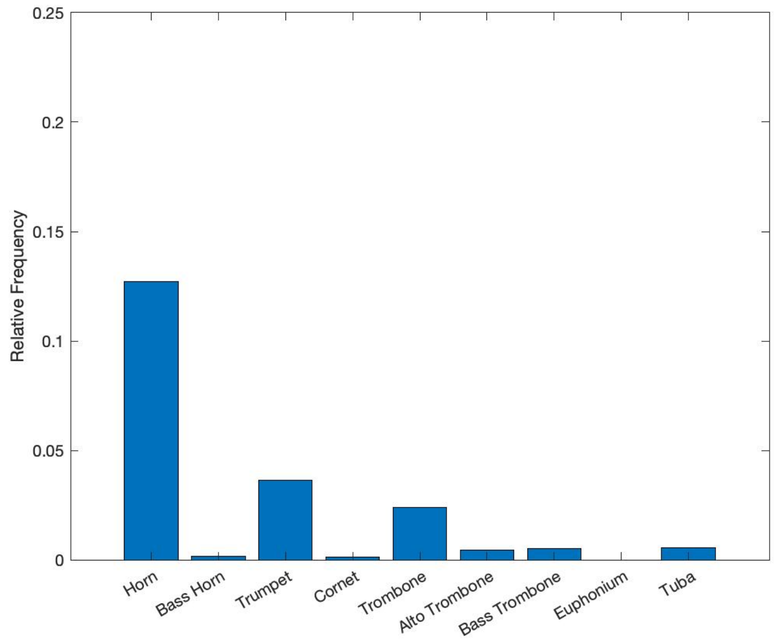

Figure 2.10 Distribution of Brass Instruments in Embellishing Instruments

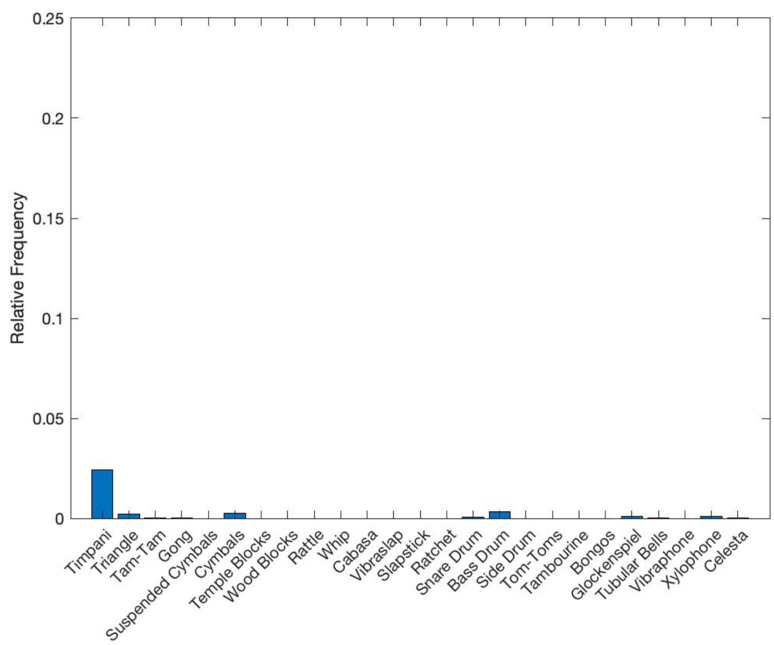

Figure 2.11 Distribution of Percussion Instruments in Embellishing Instruments

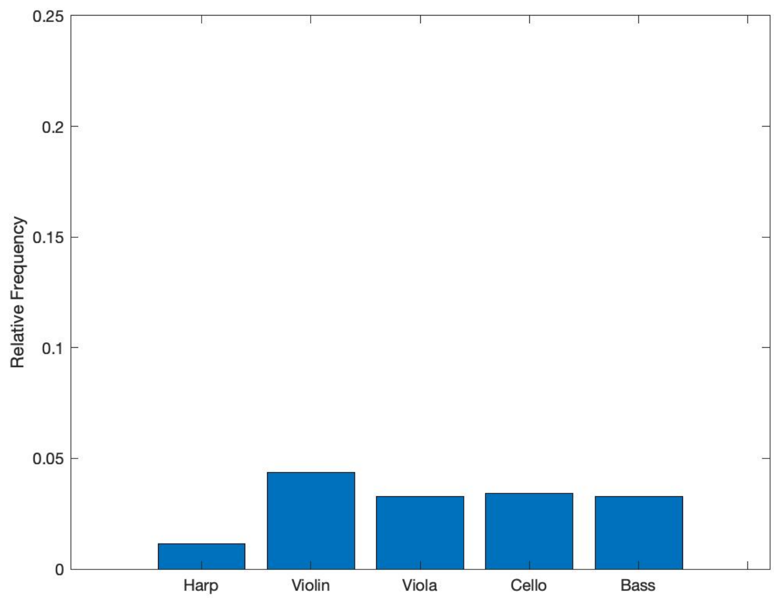

Figure 2.12 Distribution of String Instruments in Embellishing Instruments

Figure 2.8 presents the first of many differences between the embellishing instruments and the dominating instruments. The woodwinds, which had a small proportion in Figure 2.3, now form the majority of embellishing instruments at 60.43% of instances. Their distribution graph further explains this sudden change between the two blend roles. Overall, the instruments' relative frequencies have increased and the largest contributors have also changed. Although the flutes and oboes still have a significant presence, the clarinets and bassoons are the most frequently used embellishing instruments in the woodwind family. Clarinets accounted for 16% of embellishing instruments, followed closely by the bassoons at 15%. In addition, strings have a much smaller role as embellishing instruments. The violin's relative frequency has sharply dropped from its value in the dominating instruments graph (Figure 2.7), which was 22.63%, to 4.37%. Although it is still the most represented of its family, the violin now has a relative frequency that is much closer to the rest of the string instruments.

The brass instruments' proportion in Figure 2.8 is slightly lower than its proportion in Figure 2.3. Their instrument part distribution quadrant in this role is similar to that in the dominating instrument role. The most noticeable difference is that lower-pitched instruments, like the trombone and the tuba, achieve higher relative frequencies of occurrence in this blend role than as dominating instruments. Other changes include a slight increase in the use of percussion instruments. The percussion graph in Figure 2.11 shows that there is greater variety of instruments used in this blend role. Nevertheless, the most frequently used instrument remains the same, as percussion's increased proportion can be largely attributed to the timpani's relative frequency increase.

Comparison of dominating and embellishing instruments

Table 2.7 compares how often different major instruments are used as dominating instruments and as embellishing instruments in the timbral augmentation blend annotations. These percentages were calculated using the numbers found in Tables 2.2–2.5. This means that, unlike Table 2.6, these ratios include instances in which certain instrument parts may be dominating and others may be embellishing. Minor instruments, such as the bass clarinet or the bass horn, were merged in with their sibling instruments. For instance, the clarinets encompass the clarinets, piccolo clarinets, and bass clarinets.

Table 2.7: Ratio of Case-Study Instruments Dominating versus Embellishing in a Timbral Augmentation Annotation

Instrument	Dominating : Embellishing Ratio
Flute	32.30%
Oboe	15.32%
Clarinet	7.20%
Bassoon	5.08%
French Horn	45.21%
Trumpet	79.78%
Violin	178.57%

These percentages reflect what was seen in the blend roles analysis in the previous sections. The woodwind instruments all have low ratios, indicating their high frequency of use as embellishing instruments. The flutes and the oboes are higher than the clarinet and the bassoon, which explains why they were chosen to be the case-study instruments out of all the woodwind instruments. The brass instruments range between 45% and 80% with the French horns and the trumpets, respectively. Compared to the strings, the brass family is shown to be used in both blend roles at a considerably more even rate than the other instrument families. This is why the brass family is often referred to as a versatile instrument family in this work. The violins have a ratio percentage greater than 100%. This once again highlights the greater use of string instruments, namely violins, as dominating instruments in timbral augmentations.

Emerging instruments

Figure 2.13 shows the instrument family distribution for emerging instruments. Figures 2.14-2.17 add depth by showing the instrument distribution by family. They are once again constructed similarly to the other figures in this chapter.

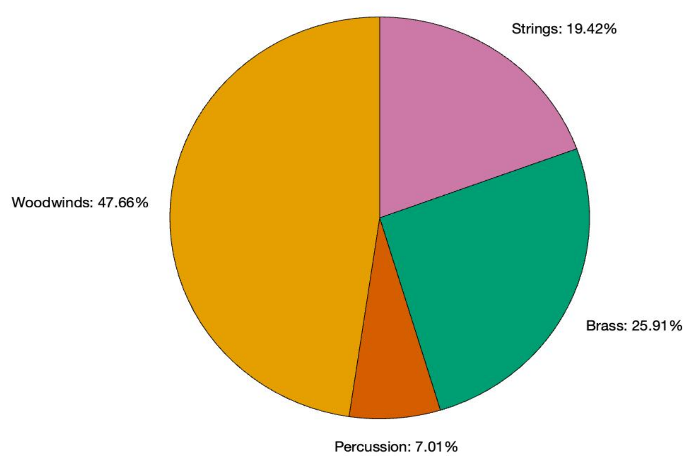

Figure 2.13 Distribution of Instrument Families in Timbral Emergence Blends

The family distribution in Figure 2.13 shares similarities with the family distribution of embellishing instruments in Figure 2.8. Both pie charts reveal that woodwinds are the most widely used instrument family in either role. In addition, the woodwind graph in Figure 2.14 shows

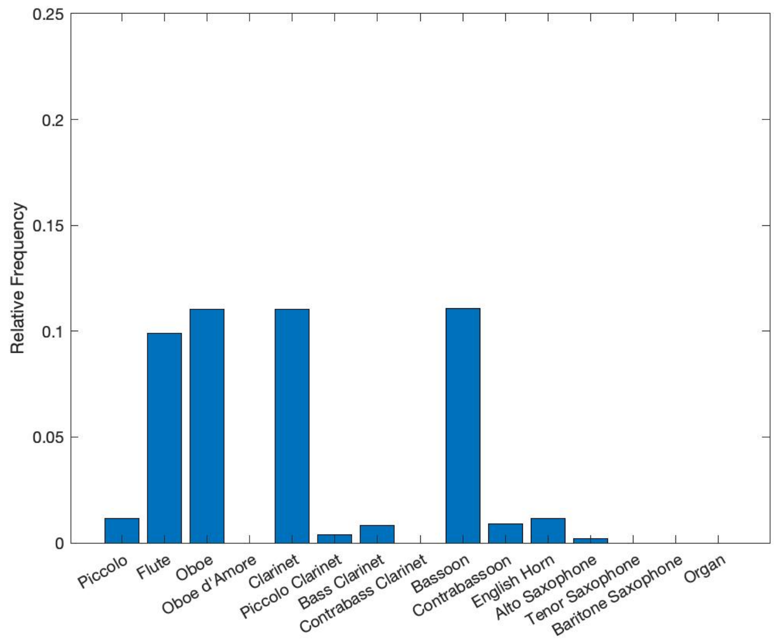

Figure 2.14 Distribution of Woodwind Instruments in Emerging Instruments

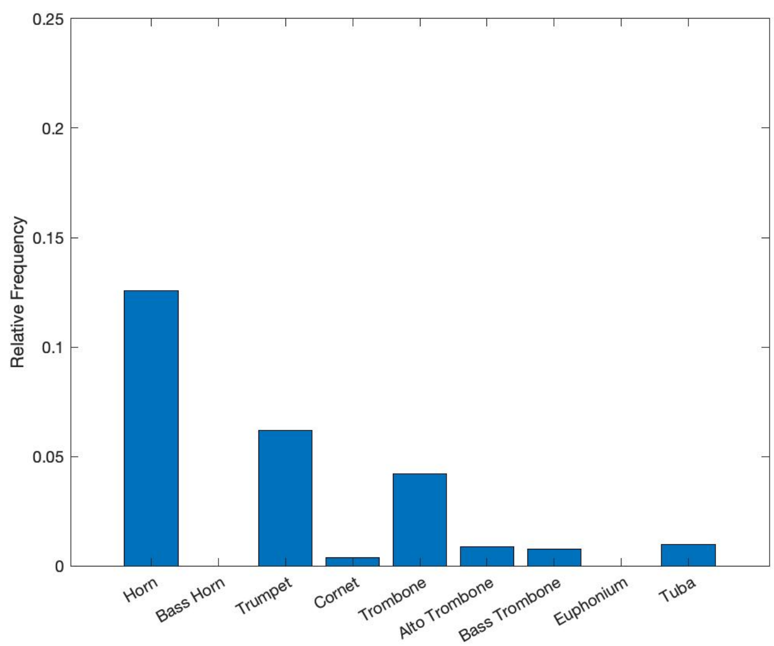

Figure 2.15 Distribution of Brass Instruments in Emerging Instruments

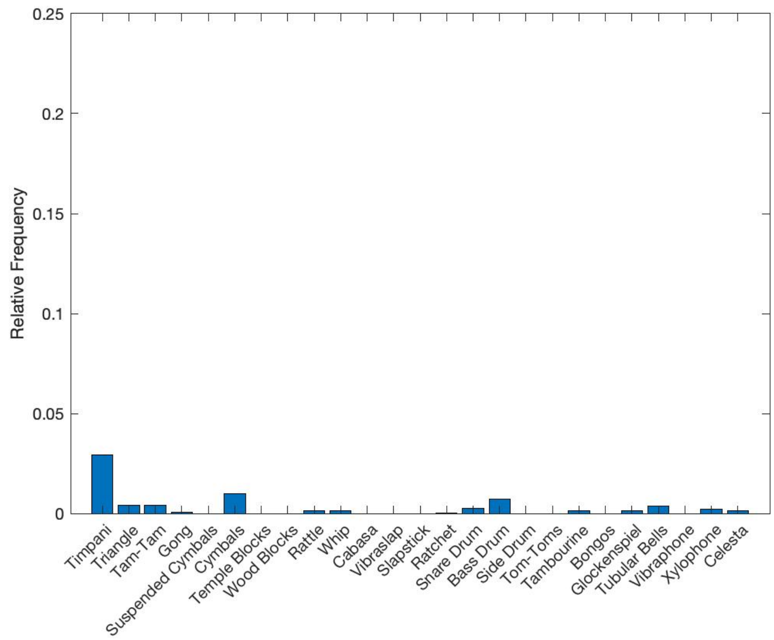

Figure 2.16 Distribution of Percussion Instruments in Emerging Instruments

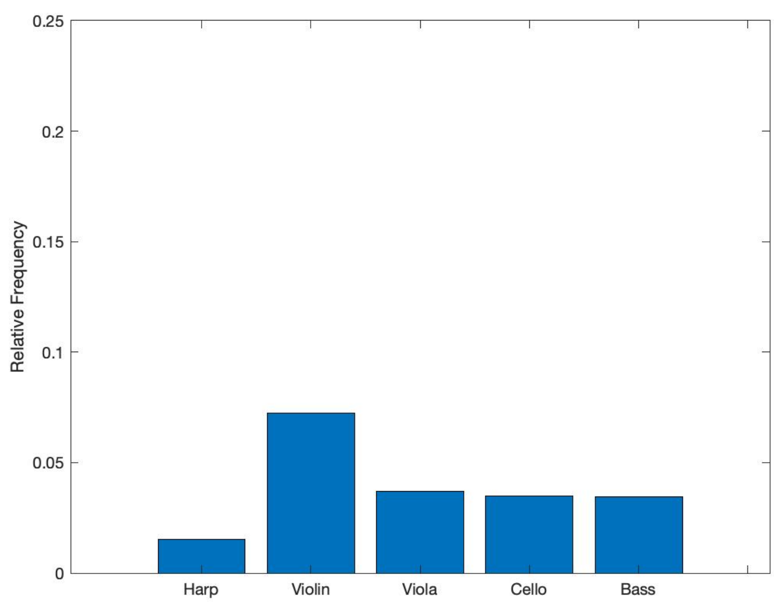

Figure 2.17 Distribution of String Instruments in Emerging Instruments

that oboes, clarinets, and bassoons are used at an approximately equal rate of 11%, with flute following closely behind at 9.9%. Once again, strings are used less frequently, especially when comparing their proportion in the emerging instrument chart (19.42%) to their proportion in the dominating instruments chart (42.70%). The string instrument distribution graph in Figure 2.17 shows that violins are still the most frequently used instruments. The instruments with lower pitch registers, including the viola, cello, and bass, have roughly equivalent relative frequencies at 3%. Brass instruments have similar proportions across all blend roles, which indicates a roughly equal presence as dominating, embellishing, and emerging instruments. Moreover, the French horn remains the most frequently used instrument in this family. The percussion family is most often used as an emerging instrument, as seen by its larger proportion in Figure 2.13. As with the other two blend roles, timpani remains the most frequently used instrument in the percussion family.

2.2.2 Analysis of case-study instruments

As previously mentioned, four instruments with high absolute counts in the dominating instruments blend role were selected as case studies. The purpose is to investigate certain instruments beyond the general blend role categories, similarly to how orchestration treatises focus on one instrument at a time. The results are reported similarly to Section 2.2.1; however, the distribution by family is omitted. Instead, the bars in each bar graph are colour-coded according to the instrument's family. Woodwind instruments are orange, brass instruments are green, percussion instruments are red, and string instruments are pink.

Violin

Figure 2.18 consists of two bar graphs showing the distribution of instruments present with a violin when it is a dominating instrument in a timbral augmentation blend. Figure 2.18a shows the instruments that are found in the same dominating group, whereas 2.18b shows the distribution of its embellishing instruments. For dominating instruments, an instance in which a single violin part is present without any other instrument parts is notated as *Violin (Solo)*. Furthermore, the presence of both violin sections is identified with the label *Violin (Multiple parts)*. This convention will be present in the other case-study instrument analyses in this chapter. Figure 2.19, on the other hand, displays the distribution of the violin's emerging instruments.

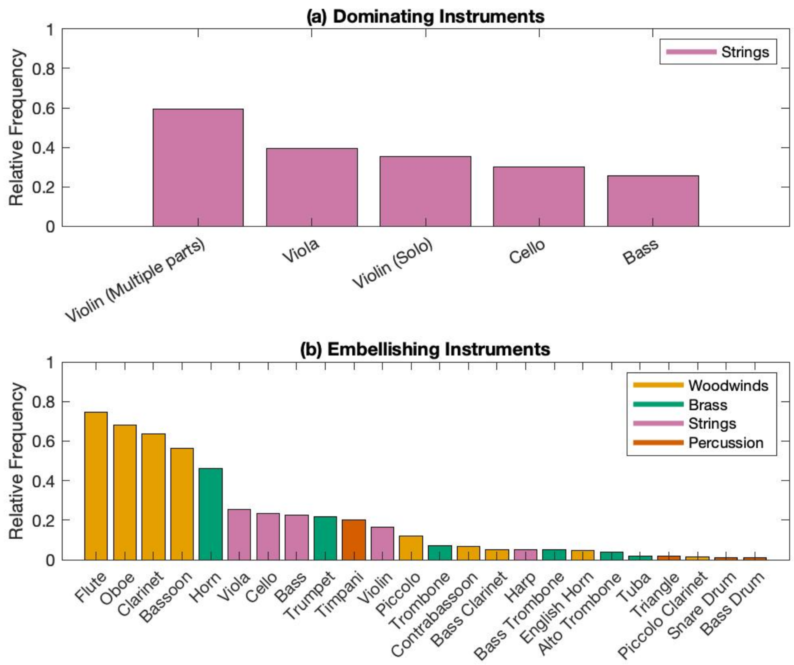

Figure 2.18 Distribution of Instruments Playing with the Violin in Timbral Augmentation Blends

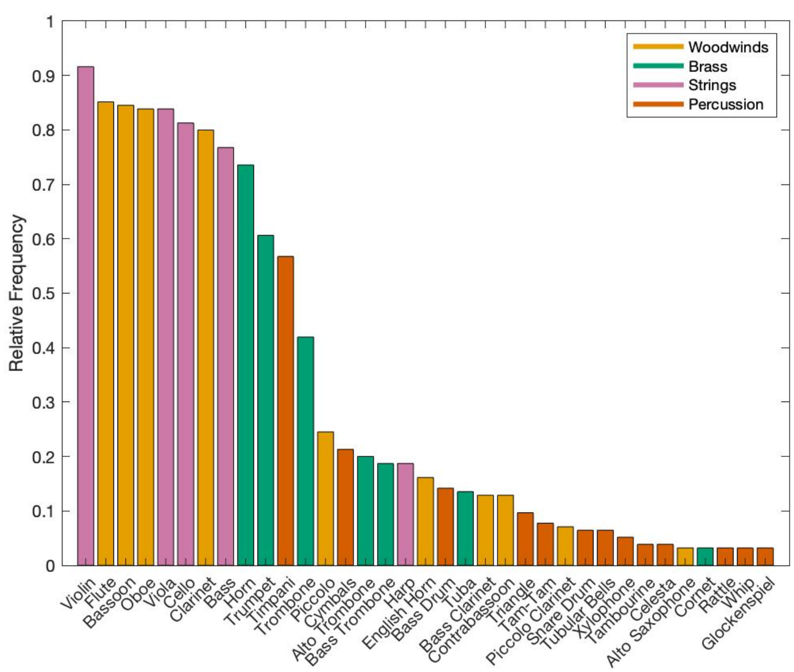

Figure 2.19 Distribution of Instruments Playing with the Violin in Timbral Emergence Blends

Figure 2.18a shows that violins largely dominate in combination with other string instruments. There were some instances of the violin dominating with a flute, clarinet, and a contrabassoon. However, these values were omitted from this graph, as they each only occurred in one annotation each. Violins are usually paired with another violin part or often found playing solo, meaning they are the only dominating instrument part in the annotation. In comparison to the handful of instruments found in 2.18a, 2.18b shows a large variety of embellishing instruments. Woodwinds are the most frequently used instrument family for embellishing the violin in a timbral augmentation blend. In particular, the flute, oboe, clarinet, and bassoon are used the most, followed by the French horn. The rest of the instruments outside these main five are all used on occasion. There is generally greater variety in the embellishing instrument blend role than in the dominating instrument blend role.

Compared to the embellishing instruments graph (in Figure 2.18b), the emerging instruments graph, shown in Figure 2.19, shows that many instruments have high relative frequencies. The highest ranked instruments appear to be a mix of timbral augmentation's top instruments for dominating instruments (i.e., violin, viola, cello, bass) and embellishing instruments (i.e., flute, bassoon, oboe, clarinet, horn) before the long tail section of less frequently used instruments begins. According to these results, both violin parts are usually found together in timbral emergence blends, occurring in 92% of these case-study annotations.

Flute

Figures 2.20 and 2.21 show the distribution of different instruments that play with the flute across the two blend types. Much like the violin case study, Figure 2.20 is a visualization of the distribution of instruments that either dominate alongside the flute or embellish it in a timbral augmentation blend, whereas Figure 2.21 shows the distribution of instruments that play with the flute during a timbral emergence blend.

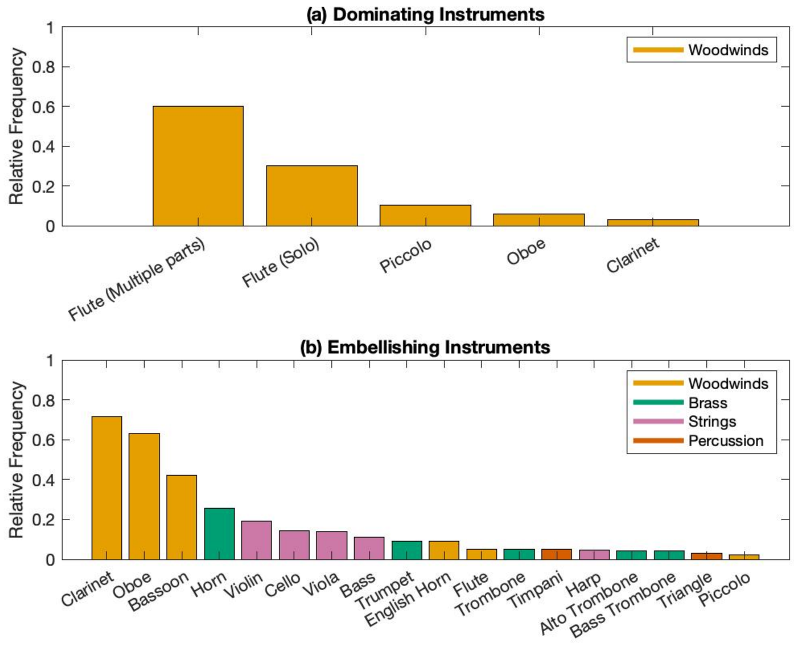

Figure 2.20 Distribution of Instruments Playing with the Flute in Timbral Augmentation Blends

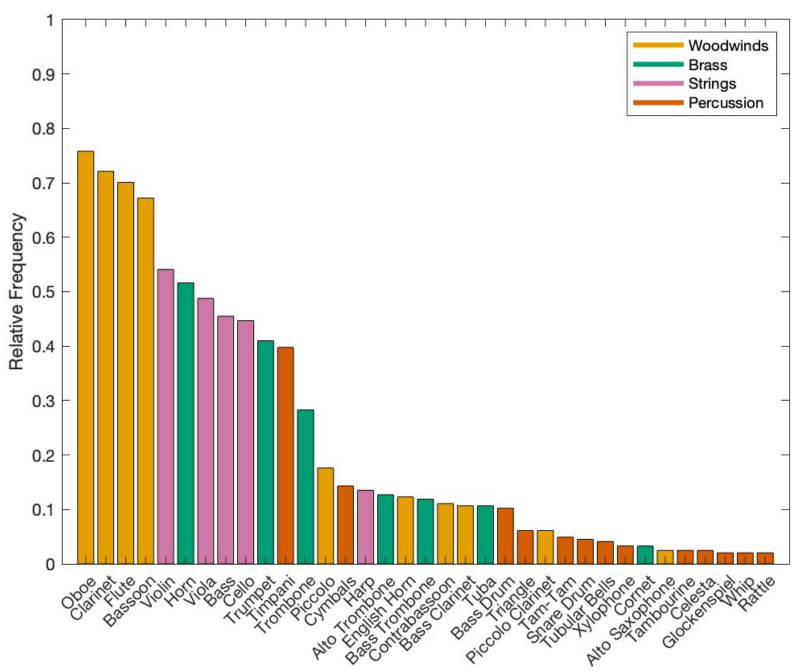

Figure 2.21 Distribution of Instruments Playing with the Flute in Timbral Emergence Blends

The timbral augmentation distribution graph (Figure 2.20) shares some features with the violin graph in Figure 2.18. 2.20a indicates that flutes tend to dominate alongside other woodwind instruments, much like the violin's tendency to dominate with other string instruments. Figure 2.20b also shows that flutes are most often embellished by woodwind instruments as well. The difference between the woodwinds in each blend role lies in the specific instruments chosen to either dominate with or embellish a flute. The highest ranked choices for the dominating instruments group in this case study are having multiple flute parts dominate, a single flute part dominating on its own, and a flute part with a piccolo part dominating conjointly. There are other woodwinds that dominate alongside the flute: the oboe and the clarinet. However, their relative frequencies of occurrence are much lower than the top three instrument options in this graph (Figure 2.20a). Oboe and clarinet achieve percentages of 5.78% and 2.89%, respectively. In addition, the annotations in which they appear are instances in which the entire woodwind section dominates the timbre at once. For instance, in the second movement of Felix Mendelssohn's *Symphony 3*, measures 205-213 feature an emergent blend of flutes, clarinets, and oboes, which become augmented by French horns and strings. Not only is this excerpt an example of a transforming blend (see section 1.3.2 in Chapter 1), but it also demonstrates how flutes can be paired with oboes and clarinets to create a woodwind dominant timbre. The distinct timbres of the three instruments are not at the forefront of this timbral augmentation blend, which is what differentiates this dominant grouping from dominant groupings of different flute parts or flutes and piccolo, for example.

Rather, the clarinet and oboe, along with the bassoon, are more often used to embellish the flute, as seen by the embellishing instruments panel. In addition to the previous observations about which instruments dominate with the flute, Figure 2.20b's relative frequencies suggest that a woodwind's timbre can influence what blend role it is assigned to in a timbral augmentation blend. The additional flute parts and the piccolo have identical or similar timbres, whereas the clarinet, oboe, and bassoon have distinctly different timbres to the flute. Woodwinds continue to be frequently used as embellishing instruments in this case study. However, these results show that the case-study instrument's timbre influences which instruments augment it in the final blend.

The most frequently used emerging instruments seen in Figure 2.21's distribution once again are a mixture of the most frequently used dominating and embellishing instruments in this case study. These top emerging instruments are followed closely by the timpani, the French horn, the trumpet, and the majority of the string family. The long tail section of less frequently used instruments in Figure 2.21 is similar to what was seen in the previous case-study instrument, and highlights the variety of instruments that can be used in a timbral emergence blend.

Oboe

Figure 2.22 shows the distribution of instruments that form timbral augmentation blends when the oboes dominate the timbre. This figure follows the same format as previous case-study figures of its kind. Figure 2.23 shows the distribution of instruments that play with the oboe in timbral emergence blends.

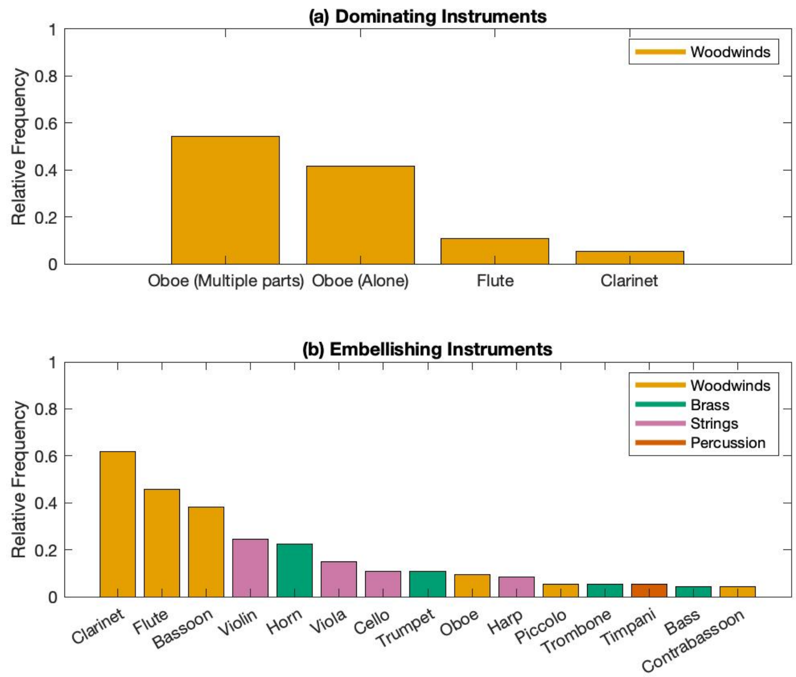

Figure 2.22 Distribution of Instruments Playing with the Oboe in Timbral Augmentation Blends

The oboe's timbral augmentation graphs share many similarities with the flute's graphs in Figure 2.20. When dominating, the oboes are most often paired with another oboe or play solo as a dominating instrument. Otherwise, oboes tend to be paired with another woodwind, but the relative frequency values of those other instruments indicate that this is a rare occurrence. Those instances are also transforming blends, meaning that the oboe and other woodwinds, like the flute or the clarinet, are not dominating the excerpt's timbre at the same time (see section 1.3.2 in Chapter 1). The relative frequencies in the embellishing instruments panel (Figure 2.22b) are lower in comparison to other case-study instruments. The clarinet and the bassoon remain

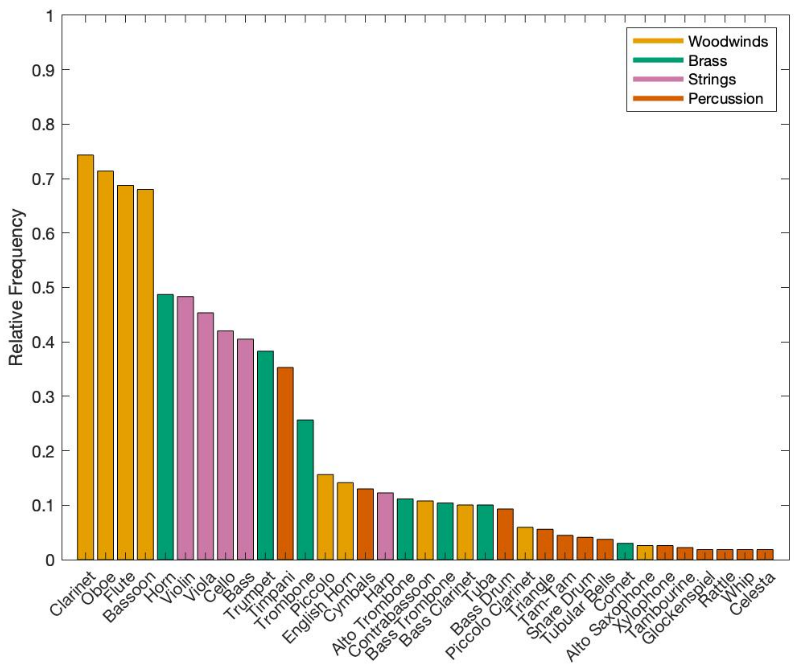

Figure 2.23 Distribution of Instruments Playing with the Oboe in Timbral Emergence Blends

some of the most frequently used embellishing instruments for the oboe, much like how they are in the flute case study.

It should be noted that there is some asymmetry with the flutes and oboes, as shown by Figures 2.20b and 2.22b. When the flute is dominating in a timbral augmentation blend, the oboe is more likely to embellish it, and vice versa. Although the notes of these annotations do not give many indicators as to why this is, such as dynamics, some of the dominating flute annotations do indicate that the flute and oboe dominating-embellishing relationship is weak and borders on timbral heterogeneity. On the other hand, the oboe and flute dominating-embellishing relationship does not have such cases. This could be due to the timbral differences between the two instruments. However, a deeper look into each score would be required to expand on this idea.

The emerging instrument distribution in Figure 2.23 is very similar to the flute's emerging instrument distribution. There are minor differences in the order of the instruments, but the overall shape to the distribution remains the same. The section of four frequently used woodwinds is roughly similar, followed by similar string and brass sections before the long tail of less used instruments begins. Oboes and flutes are often found together in timbral emergence blends, as seen by their high relative frequencies in their distribution graphs. So, it follows that their distribution graphs are very similar to each other.

French horn

Figure 2.24 shows the distribution of instruments that play with the French horn in timbral augmentation blends. The top panel, once again, shows the distribution of instruments that dominate the timbre simultaneously with the horn, whereas the bottom panel provides details on the instruments that embellish it. Figure 2.25 then displays the distribution for the French horn's emerging instruments.

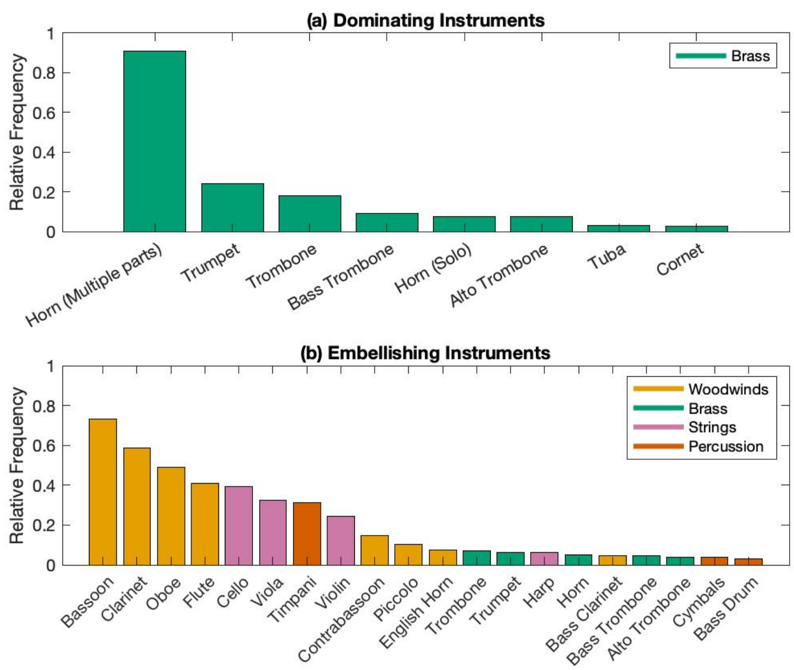

Figure 2.24 Distribution of Instruments Playing with the with French Horn in Timbral Augmentation Blends

The dominating instruments distribution once again shows that French horns are paired with other dominating instruments within their own family. However, the largest difference between this case study and the previous ones is the fact that the horn is rarely heard alone. In fact, the horn dominates with at least one additional horn part in 91% of this case-study's timbral augmentation annotations. That is, in 91% of the French horn case-study annotations, it is

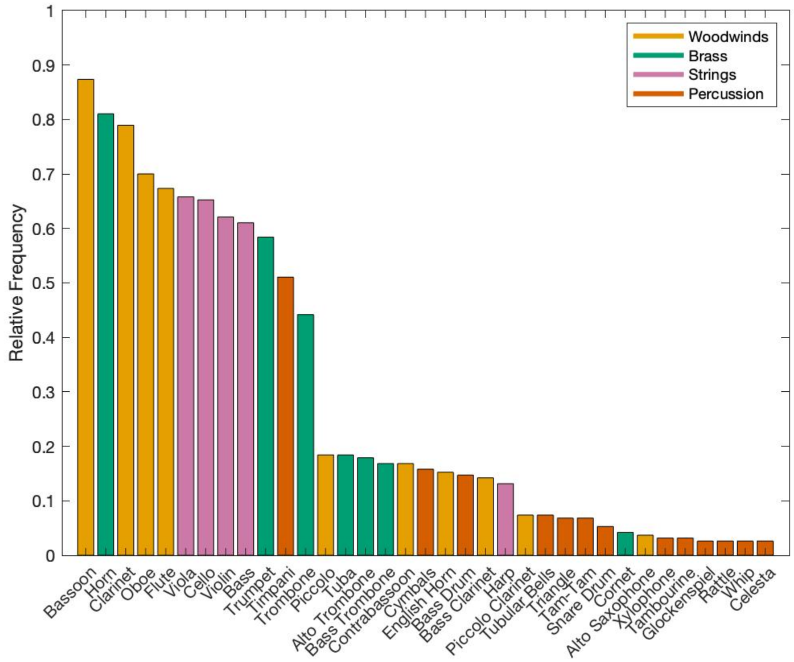

Figure 2.25 Distribution of Instruments Playing with the French Horn in Timbral Emergence Blends

accompanied by at least one other French horn part. This is the highest value seen of any case study's dominating instrument distribution graph (Figures 2.18a, 2.20a, 2.22a). Figure 2.24a indicates that horns rarely dominate as one singular horn section, as the *Horn (Solo)* value is ranked below *Horn (Multiple Parts)*. In fact, the French horn is doubled so often that other brass instruments rank higher than *Horn (Solo)*, including the trumpet, trombone, and bass trombone.

The embellishing instruments panel (Figure 2.24b) shows some similarities with what has been seen previously. Once again, woodwinds are strong choices to embellish the French horn. However, the bassoon now has the highest relative frequency of occurrence, at 73.25%, over the clarinet and the flute. In addition, the cello is now the most frequently used string instrument to embellish the horn. The distribution of emerging instruments in Figure 2.25 is once again reflective of previous emerging instrument graphs. The slight changes to the timbral augmentation's instrumentation distribution graphs are reflected in this graph as well. Bassoons are once again the most frequently used emerging instruments, at 87.27%, followed by the other three frequently seen woodwinds (i.e., clarinet, oboe, flute). The viola and cello, however, are now the most frequently used emerging string instruments instead of the violin.

2.3 Discussion

This analysis sought to detect some initial patterns in the instrumentation of different timbral blends. The results in the overall role analyses and the case study instruments accomplished just that. The most noticeable pattern to emerge is that of dominating instrument pairings consisting of instruments within the same family. This is supported by the results of the four case-study instruments. Orchestration treatises highlight this same phenomenon. Adler (1982) mentions on numerous occasions that many powerful sounds in orchestral music are the result of doubling. Adler's use of the term "power" appears to be interchangeable with the music perception term "dominating", as dominating instruments must be powerful enough, as well as timbrally distinct enough, to be heard over the other embellishing instruments in timbral augmentation blends. There are also certain instruments that are more likely to dominate in a timbral blend than others. Violins, and strings, in general, are most often used as dominating instruments according to these analyses. String sections in Classical pieces were often assigned the leading melodic line and described as the "prima donnas" of the orchestra (Read, 1979). This could explain the strings' prominence as dominating instruments and why they are often paired with other dominating string instruments for a dominant "string" timbre.

The results of this analysis also highlight woodwind instruments as a popular choice for embellishing instruments in timbral augmentations and for timbral emergence blends. This is evident in both the role analysis and the case-study findings. Many authors indirectly address this in their treatises. Adler (1982) specifically mentions that a combination of the oboe, clarinet, and flute is of great use for providing colour to a melody. These three instruments often have the highest relative frequencies of occurrence in embellishing and emerging roles, which demand a greater blend efficacy. Woodwinds also find themselves as blended instruments as there are too many weak spots in their registers for them to effectively dominate a timbral blend, as mentioned by Piston (1955) and Read (1979). These two treatises also mention that woodwinds can be easily overpowered by a few string instruments. This may explain why they are more often used as embellishing instruments or emerging instruments.

However, string instruments are not strictly dominating instruments. The oboe and flute case-study graphs show that both instruments have been embellished by strings in timbral augmentations or featured in the same timbral emergence blend. Upon investigating annotations in which strings are embellishing or emerging, it was discovered that some of these instances are when the strings are playing pizzicato. Some examples include an excerpt from Vaughan Williams' *Symphony 8*, in which piccolo 1 is augmented by flute, glockenspiel, bassoons, and the entire string section, playing pizzicato. Debussy's *La Mer* also features a timbral emergence blend in which the bassoons blend with the pizzicato cellos. In both cases, the once overpowering string instruments suddenly blend with or are embellishing woodwind instruments. These examples highlight the role that performance techniques play in orchestral blends.

The emerging instrument case-study results of the flute were surprisingly similar to the results of the oboe. The two instruments have different pitch ranges and highly contrasting tone qualities, which lead to differences in their timbral augmentation instrument distributions. However, their timbral emergence distributions show that the flutes and the oboes were often paired with the same instruments. Orchestration treatises mention that the flute and the oboe are often paired together. The two instruments are said to equalize each other, as the nasality of the oboe's timbre is neutralized by the lighter sound of the flute (Adler, 1982). The purpose of pairing the flute and the oboe together is to negate the distinctive qualities of their respective timbres. This could be why they do not dominate together, as that would require the two instrument timbres to be heard distinctly at the same time. However, Figures 2.21 and 2.23 show that the flute and the oboe are often found together in timbral emergence blends. If they are present in many of the same timbral emergence annotations, it could explain why their distribution results have a few similarities.

The findings of this analysis position brass instruments as the most versatile instrument family. That is, they have achieved similar relative frequencies of occurrence across all three blend roles. This is supported by orchestration treatises. Piston (1955) calls attention to the French horn's versatility in particular. He explains that the tone quality of the horn allows for it to be placed alongside any instrument without their timbres conflicting with each other. This adaptability among brass instruments is reinforced by the instrument family's extensive pitch range. Trumpets and French horns cover the higher pitch registers, whereas the tuba and the trombone can be used for melodic lines with lower pitches. In addition to the balanced strength of their range, especially when compared to the weak spots found in woodwinds, brass instruments can adapt to fit multiple roles in blends.

Piston's treatise also emphasizes how often the horns are doubled. The word "doubling" is often found with "power" in the instrument's description throughout his treatise. Often, the doubling done is to maintain the volume of the French horns. This translates to the French horn's dominating instrument pairing results in Figure 2.22a. An annotation in which the horn is a dominating instrument almost always has more than one horn part present. Piston adds that the melody in an orchestral piece is often interchanged between strings and brass. Since these compositions often favour the string section (Read, 1979), it follows that the French horn's doubling is done to make it as strong as the string section. Adler (1982) also notes that bassoons are often paired with French horns. He uses the term "doubling", which has been associated with instruments in a dominating instruments group, to describe how the bassoon interacts with the horn. However, our findings indicate that bassoons are used more often as an embellishment or emerging instrument for horns. This contradictory message could be explained by the difference in their instrument families. Although the bassoon may be doubling French horn's melodically, the aforementioned weaknesses in the woodwind's range could mean that it falls into an embellishing role in a timbral augmentation blend with French horn.

One of the objectives of this analysis was to identify key differences between embellishing instruments and emerging instruments. At first glance, both instrument roles demand that the instruments involved blend effectively with each other. However, the most apparent difference is their average group size. As mentioned in the basic statistics of the database (section 2.1.1), timbral augmentations involve approximately two dominating instrument parts and an average of roughly seven instrument parts to embellish the blend. Timbral emergences, on the other hand, involve an average of 13 instrument parts. This result suggests that timbral emergence blends often occur when the entire orchestra is playing together, or *in tutti*. This could also explain why so many instrument parts and types are unique to or more common in timbral emergence blends. For example, percussion instruments are more frequently used in timbral emergences as opposed to timbral augmentations. Furthermore, as seen in Table 2.2, oboes are used more frequently in timbral emergence blends than as embellishing instruments in timbral augmentations. Berlioz (1948) notes that the oboe is most often used when the orchestra plays *in tutti*, as the other instruments can offset the oboe's nasal tone quality and make its timbre less distinct. As such, although embellishing instruments and emerging instruments may have similar objectives in theory, the circumstances of their respective blend types impact their instrumentation.

It should be noted that the timbral emergence results could be influenced by the differences between punctuated and sustained timbral emergence blends (see section 1.3.2). There are a total of 108 punctuated annotations, with an average of 20.74 instrument parts involved, and 205 sustained annotations, which have an average of 9.12 instrument parts involved. These different averages show the great variety in the number of instrument parts that can be involved in the creation of a timbral emergence blend. After investigating the distribution of the relative frequencies of occurrence among the instruments, the largest difference between the two blend durations is that punctuated blends employ percussion instruments more frequently than sustained blends. This makes sense because punctuated blends have a very short duration, which makes impulsive instruments with shorter attack times a good option to include. For example, Mahler's *Symphony 1* ends with an orchestral tutti, which includes the timpani, triangle, tam-tam, cymbals, and bass drum. The result is a punctuated timbral emergence blend with many instruments involved. In comparison, earlier in *Symphony 1*, Mahler achieves a sustained timbral emergence blend using only flutes, oboes, and clarinets. The annotator refers to this excerpt as a "perfect sonority of winds".

By examining the relative frequencies of occurrence of instruments found in timbral augmentation blends and timbral emergence blends, some initial patterns and characteristics of OrchARD's blend data were revealed. This analysis also highlighted some key differences between embellishing instruments in a timbral augmentation and instruments involved in a timbral emergence. However, these relative frequencies do not answer the larger questions of the current study about the instrumentation of blends. Are there bigger patterns that cannot be detected by pairings? Does blend strength factor into what instrumentation patterns are frequently seen? These questions will be addressed in the subsequent analysis in Chapter 3.

Chapter 3

Experiment 2: Frequent Pattern Mining Using FP-Growth

The previous chapter conducted a basic investigation into the contents of OrchARD and their orchestration patterns. From that initial analysis, some observations were made that reflected the teachings outlined in orchestration treatises. However, the patterns and observed instrument groupings were restricted to pairs. For instance, the violin case-study data showed that violins are frequently augmented by flutes. However, the analysis did not show if the violins are often augmented by flutes and clarinets simultaneously. These scenarios could be counted manually, of course, but a more automated process is preferred due to the many possible instrument part combinations. In addition, the results from that analysis painted a restricted picture. Are the instrument parts identified in the frequently occurring instrument groupings statistically dependent on one another? For example, is clarinet 2's appearance in a blend role dependent on clarinet 1 or bassoon 1's presence? There also remains the question of blend strength's impact on these results. Are there some instrument groupings that occur frequently, yet consistently produce weak blends?

These queries will be addressed using frequent pattern mining. This technique was initially developed as a marketing analysis tool, but has since been adapted to numerous fields. Frequent pattern mining provides researchers with a set of tools to investigate frequently occurring items and factors that contribute to their presence in the dataset (Han et al., 2007). This makes it an interesting tool for answering the first experiment's follow-up questions.

This chapter begins with an overview of frequent pattern mining terminology and algorithms. Then, we elaborate on the methodology and results of the frequent pattern mining experiment on OrchARD's timbral augmentation and timbral emergence data.

3.1 Background on frequent pattern mining

The Han et al. (2007) review of frequent pattern mining provides definitions for each term in this field and will be paraphrased throughout this section. The dataset being mined is made up of entries, known as *transactions*. Each transaction consists of *items*. An *itemset* is a subset of items whose sizes range from a single item to an entire transaction's items. For instance, the transaction [(a), (b), (c), (d)] includes itemsets such as [(c)], [(a), (b)], or [(a), (b), (c), (d)]. Each itemset has a *support* value, which is the rate at which it occurs in the dataset. In other words, it is the itemset's relative frequency of occurrence within the transaction list. Minimum support is often used as a threshold value for determining frequent patterns in mining algorithms. If an itemset's support falls below the minimum support value, it is not considered to be a frequently occurring itemset.

The original intent of frequent pattern mining was to find *association rules*, which are the relationships between individual items or between different itemsets. They are written in the form $A \rightarrow C$, where A is the antecedent and C is the consequent. Although frequent pattern mining research interests are primarily in algorithm applications in fields other than marketing and improvements to existing methodologies, association rule mining is still discussed and researched today (Han et al., 2007). The strength of these relationships can be measured by different metrics, which are often reported alongside the frequent itemsets and association rules. For the current study, five different metrics will be used: *support*, *confidence*, *lift*, *leverage*, and *conviction*.

The *support* for $A \rightarrow C$ is the support of the combined itemset $A \cup C$ (Eq. 3.1). The latter four metrics are dependent on the itemset's support. *Confidence* is the conditional probability of the consequent appearing in a transaction given that it also contains the antecedent (Eq. 3.2). *Lift* measures how much more frequently the antecedent and consequent of rule $A \rightarrow C$ occur together than would be expected if they were statistically independent (Eq. 3.3) (Prajapati et al., 2017). *Leverage* measures the difference between the observed frequency of the antecedent and consequent appearing together and the frequency that would be expected if they were independent (Eq. 3.4). Thus, if A and C are independent, the leverage would be 0, and a negative value would suggest a negative correlation between the two values (Geng & Hamilton, 2006). *Conviction* measures the implication strength of the association rule based on the consequent's support and the rule's confidence. When A and C are independent, the conviction will equal 1. A value less than one occurs when the consequent's support is greater than its confidence, which usually occurs when the consequent occurs in nearly all transactions. In that case, the consequent does not depend on the antecedent, but the antecedent likely depends on the consequent (Prajapati et al., 2017). The equations for these four metrics are found below:

$$support(A \rightarrow C) = support(A \cup C) \qquad [0,1] \tag{3.1}$$

$$confidence(A \rightarrow C) = \frac{support(A \rightarrow C)}{support(A)} \qquad [0,1] \tag{3.2}$$

$$lift(A \rightarrow C) = \frac{confidence(A \rightarrow C)}{support(C)} \qquad [0,\infty] \tag{3.3}$$

$$leverage(A \rightarrow C) = support(A \rightarrow C) - support(A) * support(C) \qquad [-1,1] \tag{3.4}$$

$$conviction(A \rightarrow C) = \frac{1 - support(C)}{1 - confidence(A \rightarrow C)} \qquad [0,\infty] \tag{3.5}$$

3.1.1 Algorithms

The earliest and most influential frequent pattern mining algorithm is the *Apriori algorithm*. According to Han et al. (2007) and Agrawal and Srikant (1994), the Apriori algorithm is built on the principle that a frequently occurring itemset of length k must consist of k frequently occurring items. The algorithm begins by creating itemsets of length 1. Each itemset consists of each individual item featured in the dataset. The itemsets' minimum support values are calculated, and if they fall below the user-defined minimum support, they are removed. The remaining itemsets of length one are then paired to form a new set of itemsets, now of length 2. This process repeats, each time producing itemsets of length $k + 1$ from the previous set of k-length itemsets. Once the algorithm can no longer produce any new itemsets, it terminates, leaving the final frequent patterns (Aguilar-Ruiz et al., 2013). Figure 3.1 shows the pseudocode for the Apriori algorithm.

Algorithm 1 Apriori

$C_1 = \mathcal{A}(\mathbb{X})$ is the set of all one-itemsets, $k = 1$
while $C_k \neq \emptyset$ **do**
  scan database to determine support of all a_y with $y \in C_k$
  extract frequent itemsets from C_k into L_k
  generate C_{k+1}
  $k := k + 1$.
end while

Figure 3.1 The Apriori algorithm pseudocode. Taken from Hegland (2007)

Although the Apriori algorithm was revolutionary at the time of its creation, it has a large flaw. Han et al. (2007) explain that, by design, the algorithm must scan the entire dataset at every iteration. This does not pose a problem for smaller datasets as the runtime difference between

Apriori's method and a more strategic approach is hardly noticeable. However, this inefficiency becomes more evident in larger datasets as it drastically increases the algorithm runtime and its computational complexity (Heaton, 2016). The Apriori algorithm's runtime issue led to different researchers developing new mining algorithms to both address Apriori's issues and to introduce new features.

FP-Growth algorithm

The *Frequent Pattern Growth algorithm* (FP-Growth) was developed to circumvent the repeated scanning of the dataset when mining for frequent itemsets. Rather than keeping the transactions on a horizontal, data is stored vertically in a tree structure. Han et al. (2007)'s review, which is co-written by one of the original creators of the FP-Growth algorithm (Han et al., 2000), explains that this method allows the algorithm to approach frequent pattern mining through a "divide-and-conquer" approach. In other words, the algorithm is able to work recursively by breaking down a problem into sub-problems, then recombining those smaller solutions to resolve the original problem. This is a common technique in computer science algorithms to reduce a problem's computational runtime (Cormen et al., 2022).

Before the FP-Growth algorithm begins, the dataset is sorted so that each transaction has its items organized in decreasing order of occurrence. These transactions are then placed into a *Frequent Pattern Tree* (FP-Tree) as branches. If a transaction and an existing branch share a *prefix*, meaning that they begin with the same itemset, then the transaction will be added to the end of that prefix as a new branch. Each tree node has a counter variable to indicate how frequently the item appears within the context of its prefix branch. A header table is also used to keep a record of each item's total number of occurrences, as well as a linked list of each item's placement in the FP-Tree (Grahne & Zhu, 2003).

Figure 3.2 provides an example of a dataset and its transactions represented as an FP-Tree. The transaction list is found in column (a). Each item is counted, and those counts are recorded in the header table in column (b). The ordering of items with the same counts, like e, c, and a, is of less importance. Although it may change how the final FP-Tree looks, it has no impact on the generated frequent itemsets when the FP-Growth algorithm is applied. Each transaction is sorted according to the header list ranking, and then placed into the tree as a tree node. Each node includes the item, its count, and a pointer connecting it to its location in the header table. If the item occurs in different branches, the node has a pointer going to its next location in the FP-Tree. For instance, $[a, b, c, e, f, o]$ becomes $[e, c, a, b, f, o]$, which can be found in the branch to the left of the root node. In this case, items that occur only once, like o, were omitted from the header table and the FP-Tree. The transaction $[a, c, d, e, g]$, when sorted, becomes $[e, c, a,$

g, d]. It shares the prefix [e, c, a] with the left branch. So, the transaction gets added to the left branch and each shared item gets its count increased by 1. The different items, g and d, are placed in a new branch created at the a node. The transaction [e, i] increases the e node in the left branch by 1. No new branch is created, as i only appears once in the dataset, so it is omitted. In comparison, the transaction [a, c, g] becomes [c, a, g] when sorted. Since it does not share a prefix with the left branch, it is added as an entirely new branch on the right of the FP-Tree's root. However, the two branches do have some common items. A pointer is thus created from the left branch's nodes to the newly added nodes in the right branch. The process repeats for the other transactions in this list until they are all added to the FP-Tree.

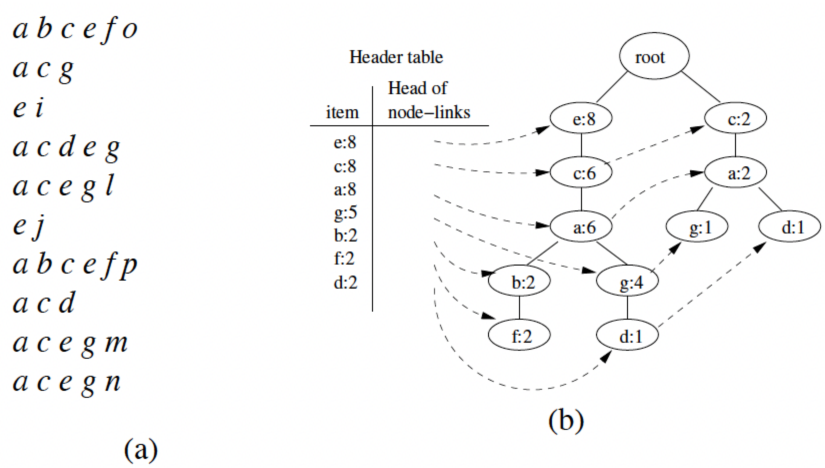

Figure 3.2 Diagram of an example dataset and its FP-Tree representation. Taken from Grahne and Zhu (2003)

Once the FP-Tree is built, FP-Growth begins with the least frequently occurring item, and uses the header table to find all instances of that item within the FP-Tree. It should be noted that this least frequently occurring item in the FP-Tree must still meet the minimum support requirement. Paths from this item to the root are stored and used to build the item's *conditional FP-Tree* (Nalousi et al., 2021). Items in the path are only added to the conditional FP-Tree if they are greater than the minimum support. Frequent patterns are then mined from the paths of this smaller tree structure. This is repeated for all items in the header table that are above the minimum support threshold, and the final set of frequent patterns across the conditional FP-Trees are the frequent patterns for the dataset (Grahne & Zhu, 2003; Han et al., 2007; Nalousi et al., 2021). Figure 3.3 outlines this algorithm in pseudocode.

Returning to our example in Figure 3.2, the FP-Growth algorithm would begin with d. Assume that the user-defined minimum support for this algorithm is 0.50, which means that the itemsets

Algorithm *FP-growth*(FP-Tree of frequent items: $\mathcal{FPT}$, Minimum Support: *minsup*,
Current Suffix: P)
begin
if $\mathcal{FPT}$ is a single path
then determine all combinations C of nodes on the
path, and report $C \cup P$ as frequent;
else (Case when $\mathcal{FPT}$ is not a single path)
for each item i in $\mathcal{FPT}$ **do begin**
report itemset $P_i = \{i\} \cup P$ as frequent;
Use pointers to extract conditional prefix paths
from $\mathcal{FPT}$ containing item i;
Readjust counts of prefix paths and remove i;
Remove infrequent items from prefix paths and reconstruct
conditional FP-Tree $\mathcal{FPT}_i$;
if ($\mathcal{FPT}_i \neq \phi$) **then** *FP-growth*($\mathcal{FPT}_i, minsup, P_i$);
end
end

Figure 3.3 Pseudocode for the FP-Growth algorithm, taken from Aggarwal (2015)

must appear at least 5 times. d, f, and b are not considered based on this threshold. g occurs 5 times, so FP-Growth begins with that item. It appears in two separate branches, so the two paths from the root to g are extracted: [a, c, e], which occurs 4 times, and [a, c], which occurs once. The conditional FP-Tree for g is then built from these paths, which are known as g's *conditional pattern base* (Han et al., 2000). The resulting conditional FP-Tree only has one branch from the root, which consists of a and c. Both items occur 5 times each. Since e only occurs 4 times, it is not included in the conditional FP-Tree. So, the only frequent pattern produced from this iteration of the algorithm is [g, a, c]. This process is repeated for the other items. The final frequent patterns produced are [g, a, c], [a, c], [a, c, e], and [c, e].

The results of our example have some itemsets that are subsets of a larger itemset, such as [g, a, c] and [a, c]. FP-Growth does not generate *maximal frequent patterns*, which means that the frequent itemsets reported by the algorithm may include some subsets of other frequent itemsets. Other algorithms, like FP-Max, have been created to account for this, as the generation of additional items may become an issue for larger datasets, especially when the minimum support value is set low (Han et al., 2007).

3.2 Method

As discussed at the beginning of this chapter, the objective of this frequent pattern mining analysis was to expand on the previous chapter's findings. Although insightful, those initial results were restricted to pairings and left many questions unanswered regarding instrumentation patterns in timbral augmentation blends. This made frequent pattern mining an ideal tool for conducting an in-depth analysis. Although the frequent patterns were expected to be similar to those found in the previous analysis, the association rule metrics were expected to provide clarity to the relationships between the instrument parts. In addition, those patterns may be expanded to include other frequently occurring instruments. Similarly to the previous chapter, OrchARD's timbral emergence blend data were examined in addition to the timbral augmentation blends. The case-study instruments (violin, flute, oboe, French horn) presented in the relative frequency analysis were also revisited.

This analysis also introduced blend strength as a variable to the pattern mining. McAdams et al. (2022) emphasized that not all blends are created equal. The degree, or strength, of a blend depends on multiple variables, including the orchestration techniques employed, the room's acoustics, and the recording techniques. This opened questions regarding the influence that blend strength might have on the resulting frequent patterns. Are there some instrumentation patterns that frequently occur strictly at lower degrees of blend? If we remove the annotations at those strengths, is there an impact on the reported frequent itemsets? Our hypothesis was that there would be some patterns that occur exclusively at weaker blend strengths. So, by including an additional lower bound that removes those weaker blends, the generated frequent patterns would change.

3.2.1 Algorithm Implementation

The frequent pattern mining algorithm used was a variation of the FP-Growth algorithm. This algorithm was chosen for its aforementioned computational efficiency and data structure. While OrchARD is not a large dataset, it is continuously growing. In order to preserve this implementation for future research, FP-Growth was selected instead of the Apriori algorithm, as it is recommended for use in larger datasets (Heaton, 2016).

Frequent pattern mining algorithms always leave the selection of minimum support at the user's discretion. This implies that the user is entering the analysis with an idea of their database's contents in order to select a good support threshold, lest they find that value through trial and error (Zhang et al., 2008). Frequent pattern mining is generally interested in the most frequently occurring results. A large number of results can be generated if the support is set low enough. However, the current study is interested in what occurs the most frequently. As such, the minimum support value was adjusted for each blend role and case-study instrument to avoid over-generating results that are not of great interest to this analysis. The values were set low enough to generate some results, and then readjusted to prevent too many results from appearing. As will be seen in the Results section, this led to different minimum support values across the different blend roles and different case-study instruments.

The structure of the FP-Growth algorithm presented in Figure 3.3 needed to be adjusted to include blend strength as an additional user-defined input. Although the minimum support value could be assigned any decimal value between 0 and 1 inclusive, the minimum blend strength was defined as a natural number between 1 and 5. During the construction of the FP-Tree, each transaction's blend strength was measured. If it fell below the chosen minimum blend strength, the transaction was discarded. This was done to remove the additional computation time of building a tree only to remove many branches after the fact, especially if the minimum blend strength was set to a high value. Once that was done, the FP-Growth algorithm proceeded as expected. Minimum blend strength was consistently set to 2 to generate the reported results. However, this value was toggled to observe its effect on the results. Following the execution of the FP-Growth algorithm, the results were presented along with the association rule metrics. As mentioned in section 3.1, the selected metrics were support, confidence, lift, leverage, and conviction.[1]

3.3 Results

3.3.1 Analysis of blend roles

Table 3.1 shows the generated frequent itemsets and the values of their association rule metrics in the three different blend roles. The itemsets column shows all frequent itemsets identified by the FP-Growth algorithm. For dominating instruments, the minimum support was set to 0.20. The reasons for this are related to some of the dominating instruments characteristics previously described in Chapter 2 (see section 2.1.1), including its instrument family division and its small group size. As mentioned in the previous section (3.2.1), the support values were chosen to avoid generate too many or too few results. The embellishing instruments' minimum support was set to 0.30 for a similar reason. The emerging instruments did not encounter this problem and their minimum support was set to 0.60.

[1]This variation of the FP-Growth algorithm was implemented in a Jupyter notebook file using Python. The code is available upon request from stephen.mcadams@mcgill.ca.

Table 3.1: Association Rule Metrics for the Frequent Patterns in Blend Roles

Blend Role	Itemset	Support	Ant.	Cons.	Support Ant.	Support Cons.	Conf.	Lift	Lev.	Conv.
Dominating Instruments	violin 1	0.330	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	violin 2, violin 1	0.208	violin 2	violin 1	0.227	0.330	0.915	2.770	0.133	7.838
Embellishing Instruments	clarinet 1	0.573	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	clarinet 2, bassoon 1, clarinet 1	0.346	clarinet 2	bassoon 1, clarinet 1	0.530	0.354	0.652	1.781	0.152	1.822
	bassoon 1, clarinet 1	0.366	bassoon 1	clarinet 1	0.560	0.573	0.653	1.141	0.045	1.232
Emerging Instruments	bassoon 1, clarinet 1	0.617	bassoon 1	clarinet 1	0.740	0.760	0.833	1.096	0.054	1.440
	clarinet 1	0.760	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Note: *Ant.*: antecedent, *Cons.*: consequent, *Conf.*: confidence, *Lev.*: leverage, *Conv.*: conviction

The support values for the antecedent and consequent were also included. All frequent patterns generated by the FP-Growth algorithm appear in the results tables. However, some frequent itemsets only consisted of one item and do not have the association rule format of $A \rightarrow C$. For those cases, only their support value was reported.

Dominating instruments

The minimum support value in the dominating instruments blend role had to be set lower than other blend roles, and this is because the itemsets found in this role had low support values. For instance, [violin 1]'s support of 0.330 means that violin 1 appeared in 33% of all dominating instrument instances. The frequent itemsets, [violin 1] and [violin 1, violin 2], are both comprised of string parts. This instrument family was found to be more frequently seen as dominating instruments in Chapter 2 (section 2.2.1.1). Although other instruments, like the flute, oboe, and French horn, also had high relative frequencies of occurrence, they did not meet the minimum support, meaning that they did not occur frequently across all dominating instrument instances.

[violin 2, violin 1]'s relationship is clarified by the association rule metrics. Their joint support value is low (0.208), but their confidence is high (0.915). This means that, if violin 2 is a dominating instrument part in a timbral augmentation blend, it appears along with the violin 1 part

91.5% of the time. The lift and conviction values are greater than 1, which indicates a statistical dependence between the two instrument parts. [violin 2, violin 1]'s conviction value is the highest found in this table, showing that the two instruments are highly related to each other.

Embellishing instruments

The support values for the frequent itemsets detected in this data subset are slightly higher than those for the dominating instruments. Clarinet 1 is found in all frequent itemsets, achieving a support of 0.573. The other detected frequent itemsets are [clarinet 2, bassoon 1, clarinet 1] and [bassoon 1, clarinet 1]. All frequent items across all itemsets are woodwind instruments, specifically the two woodwind instruments that had a large increase in presence between the dominating and embellishing blend roles in Chapter 2 (see section 2.2.1.2).

The other association rule metrics indicate that these instrument parts are statistically independent of each other. The confidence values are not as high in this blend role (0.652, 0.653) as they are in the dominating instruments (0.915) and emerging instruments (0.833). The lift and conviction values are both near 1, and the leverage values are near zero. In other words, the values of these three metrics indicate that the items in the frequently occurring itemsets are statistically independent of each another (see section 3.1). The [clarinet 2, bassoon 1, clarinet 1] itemset produced higher values of leverage (0.152), although they are still not strong indicators for statistical dependence between clarinet 2 and [bassoon 1, clarinet 1]. These itemsets may occur frequently, but their items are not statistically dependent on one another.

The higher lift and conviction values in the [clarinet 2, bassoon 1, clarinet 1] itemset may be detecting that the clarinet is being doubled based on its presence in the antecedent and consequent. Recall that the [violin 2, violin 1] itemset in the dominating instruments blend role had a lift value of 2.770, a leverage value of 0.133, and a conviction value of 7.838. [clarinet 2, bassoon 1, clarinet 1]'s values are not as high, but they are still higher than the other embellishing instruments itemset. The shared trait between these two itemsets is that they both include a doubled instrument. Violin 1 is doubled in the [violin 2, violin 1] itemset, and clarinet 1 is doubled in the [clarinet 2, bassoon 1, clarinet 1] itemset. Therefore, a link can be made between the doubling of an instrument and statistical dependence.

Emerging instruments

The emerging instrument frequent itemsets have the highest average support values of all three blend roles. The detected emerging instrument itemsets, [bassoon 1, clarinet 1] and [clarinet 1], contain similar items as the itemsets found in the embellishing instrument frequent patterns. Their support values are significantly higher than those of the embellishing instruments. Clarinet

1, in particular, has the highest support value of any item across all three blend roles at 0.760. [bassoon 1, clarinet 1]'s confidence value is notably higher as an emerging instrument itemset than as an embellishing instrument itemset, implying that the likelihood of clarinet 1 appearing with bassoon 1 is greater in a timbral emergence blend. However, its relatively low lift and leverage indicate that the two instrument parts are not statistically dependent on each other. Furthermore, the low conviction value shows that the two items do not have a strong relation to one another.

It is interesting to note that, despite the flute and the bassoon having high relative frequencies in the embellishing and emerging blend roles (see sections 2.2.1.2 and 2.2.1.3), the clarinet is the only instrument that achieves a consistently high support value in Table 3.1. This cements the clarinet's importance as an embellishing and emerging instrument over the other two instruments. The bassoon does appear frequently as an embellishing and emerging instrument. However, it does not have the same versatility as the clarinet. Rather, the bassoon is used more often with brass instruments, as will be explored in the upcoming French horn analysis (see section 3.3.2.4).

3.3.2 Analysis of case-study instruments

Violin

Table 3.2 shows the resulting frequent itemsets that were detected in this case-study's data. The minimum support was set at 0.40, 0.50, and 0.70 for the dominating instruments, the embellishing instruments, and the emerging instruments, respectively.

Table 3.2: Association Rule Metrics for the Frequent Patterns in the Violin Case Study

Blend Role	Itemset	Support	Ant.	Cons.	Support Ant.	Support Cons.	Conf.	Lift	Lev.	Conv.
Dominating Instruments	violin 1	0.966	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	violin 2, violin 1	0.610	violin 2	violin 1	0.667	0.966	0.915	0.947	−0.034	0.400
Embellishing Instruments	flute 1	0.766	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	oboe 1, flute 1	0.604	oboe 1	flute 1	0.695	0.766	0.869	1.133	0.071	1.781
Emerging Instruments	flute 1, violin 2, violin 1	0.760	flute 1	violin 2, violin 1	0.880	0.927	0.894	0.965	−0.029	0.691
	violin 2, violin 1	0.927	violin 2	violin 1	0.940	1.000	0.986	0.986	−0.013	0.000
	oboe 1, flute 1, violin 2, violin 1	0.780	oboe 1	flute 1, violin 2, violin 1	0.847	0.787	0.921	1.171	0.114	2.709
	violin 1	1.000	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

The first part of every case-study instrument has a very high support value. Often, it appears in every transaction. In the violin case-study results, violin 1 has a support value of 1.000 in the emerging instruments, but has a lower value of 0.966 in the dominating instruments. This means that there are instances where violin 2 dominates without violin 1. Due to this constant presence, some association metrics are affected, as it is challenging to determine an instrument part's dependence, or lack thereof, on other instruments if that instrument part is found in every transaction. This disproportionately affects the dominating and emerging instruments, as the instances where violins embellish a timbral augmentation blend were not considered in case studies.

The resulting frequent itemsets across all three blend roles are consistent with the results from the relative frequency analysis (see section 2.2.2.1). The dominating instrument itemsets consist solely of violin parts. The two resulting frequent itemsets for the violin's embellishing instruments are [flute 1] and [oboe 1, flute 1]. The frequent itemsets found in the emerging instruments are a combination of the items that make up the dominating and embellishing frequent itemsets: [flute 1, violin 2, violin 1], [violin 2, violin 1], [oboe 1, flute 1, violin 2, violin 1], and [violin 1].

Based on the previous results, both in Table 3.1 and in the previous chapter's relative frequency analysis (see section 2.2.2.1), the items in [violin 2, violin 1] are expected to be somewhat statistically dependent on each other. The confidence value (0.915) reflects this expectation in both the dominating and emerging instruments. Its value confirms that violin 2 is not likely to appear without violin 1. However, the lift, leverage, and conviction values indicate something different. As mentioned in this chapter's background section (3.1), lift and conviction values that are less than 1 and a negative leverage value indicate a negative correlation. As such, [violin 2, violin 1]'s items are statistically independent of each other in the dominating and emerging instrument blend roles. This could be an effect of violin 1's presence across all transactions. If violin 1 appears in nearly all transactions, it follows that it has little to no statistical influence on violin 2's presence.

These observations are contrasted with the embellishing instruments results. The association metrics for [oboe 1, flute 1] are expected when compared to the results in the blend role mining. Although the confidence value is high, at 0.869, the lift, leverage, and conviction values indicate statistical independence. In other words, although flute 1 is often embellishing if oboe 1 is also embellishing, the two items are not necessarily dependent on one another. This reflects what was seen in section 3.3.1.2, in which it was discussed that different instruments are not statistically dependent on each other. The [oboe 1, flute 1, violin 2, violin 1] itemset's metrics are also in line with some of the results from the blend role analysis. This itemset achieved the highest conviction value at 2.709. Despite the presence of [violin 2, violin 1] in this itemset, it seems that the presence of oboe 1 and flute 1 in this itemset may have offset the influence of violin 1's ubiquity. [flute 1, violin 2, violin 1] also introduces another instrument to [violin 2, violin 1]. However, the consequent's support is higher than the antecedent's support, which indicates a negative correlation, as confirmed by its lift and leverage values.

Flute

Table 3.3 reports the resulting frequent itemsets and their association rule metrics for this case study. Minimum support was set to 0.50, 0.40, and 0.70 for the dominating instruments, the embellishing instruments, and the emerging instruments, respectively. The decision process for these values follows a similar technique to that used for the violin case-study and the blend role analysis (see section 3.2.1).

The frequent itemsets produced across all three blend roles are in agreement with the results from the relative frequency analysis (see section 2.2.2.2). Flutes are the only items that appear in the resulting dominating instrument frequent itemsets. Both the embellishing and emerging instrument frequent itemsets consist of clarinets and oboes, with the emerging instruments also including flutes. Flute 1 has a support value of 1.000 as both a dominating instrument and an emerging instrument. These values mean that flute 1 is present as a dominating instrument in all timbral augmentation blends examined in this case study, as well as all timbral emergence blends. Unlike violin 2 in the violin case study, flute 2 is never dominating without flute 1. The confidence and lift values are the same for [flute 2, flute 1], which results in a conviction value of 0. Nevertheless, the lift value is closer to 1 and the leverage value is closer to 0 compared to [violin 2, violin 1]'s values in Table 3.3's dominating instruments blend role.

Table 3.3: Association Rule Metrics for the Frequent Patterns in the Flute Case Study

Blend Role	Itemset	Support	Ant.	Cons.	Support Ant.	Support Cons.	Conf.	Lift	Lev.	Conv.
Dominating Instruments	flute 1	1.000	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	flute 2, flute 1	0.627	flute 2	flute 1	0.634	1.000	0.990	0.990	-0.006	0.000
Embellishing Instruments	clarinet 1	0.758	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	oboe 1, clarinet 2, clarinet 1	0.441	oboe 1	clarinet 2, clarinet 1	0.665	0.658	0.664	1.008	0.003	1.015
	clarinet 2, clarinet 1	0.658	clarinet 2	clarinet 1	0.671	0.758	0.981481	1.295	0.150	13.081
Emerging Instruments	flute 1	1.000	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	flute 2, clarinet 1, oboe 1, flute 1	0.672	flute 2	clarinet 1, oboe 1, flute 1	0.867	0.841	0.840	0.999	-0.001	0.995
	clarinet 1, oboe 1, flute 1	0.841	clarinet 1	oboe 1, flute 1	0.887	0.933	0.948	1.016	0.013	1.281
	oboe 1, flute 1	0.933	oboe 1	flute 1	0.933	1.000	1.000	1.000	0.000	inf

The embellishing instrument frequent itemsets both include clarinets, which is in line with what was reported in the flute's relative frequencies of occurrence. The confidence, lift, leverage, and conviction values are favourable for [clarinet 2, clarinet 1], suggesting that these two instruments are statistically dependent on each other. However, the [oboe 1, clarinet 2, clarinet 1] itemset shows statistical independence between oboe 1 and the two clarinet parts based on its lift, leverage, and conviction values. This, in addition to previous instances of high conviction values in Table 3.1, supports the ongoing observation that statistical dependence is greater in parts from the same instrument as opposed to parts from different instruments.

It was noted in the violin case study that additional instruments can offset the effects of having an item that is present in all transactions. That same phenomenon can be seen in this case study with the [clarinet 1, oboe 1, flute 1] itemset in the emerging instruments blend role. Although there is statistical independence between clarinet 1 and [oboe 1, flute 1], it does break the trend of the antecedent and consequent being negatively correlated, which is a result of [flute 1] having a support of 1.000. This is not the case for [flute 2, clarinet 1, oboe 1, flute 1], as the antecedent, [flute 2], doubles the omnipresent flute 1. It should also be noted that [oboe 1, flute 1] has a confidence of 1.000 due to flute 1's high support value, which in turn causes an infinite conviction value.

Oboe

Table 3.4 shows the results of the FP-Growth algorithm for the oboe case study, including the association rule metrics. The minimum support was set to 0.40, 0.50, and 0.70 for the dominating instruments, the embellishing instruments, and the emerging instruments, respectively.

Table 3.4: Association Rule Metrics for the Frequent Patterns in the Oboe Case Study

Blend Role	Itemset	Support	Ant.	Cons.	Support Ant.	Support Cons.	Conf.	Lift	Lev.	Conv.
Dominating Instruments	oboe 1	1.000	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	oboe 2, oboe 1	0.544	oboe 2	oboe 1	0.589	1.000	0.925	0.925	-0.044	0.000
Embellishing Instruments	clarinet 1	0.622	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	clarinet 2, clarinet 1	0.533	clarinet 2	clarinet 1	0.556	0.622	0.960	1.543	0.188	9.444
Emerging Instruments	flute 1, oboe 2, clarinet 1, oboe 1	0.707	flute 1	oboe 2, clarinet 1, oboe 1	0.860	0.805	0.822	1.021	0.015	1.095
	oboe 2, clarinet 1, oboe 1	0.805	oboe 2	clarinet 1, oboe 1	0.884	0.907	0.911	1.005	0.003	1.040
	clarinet 1, oboe 1	0.907	clarinet 1	oboe 1	0.913	1.000	0.995	0.995	-0.005	0.000
	oboe 1	1.000	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Much like the flute case study, the oboe's results are similar to its relative frequency analysis results (see section 2.2.2.3). The frequent itemsets detected for dominating instruments consist solely of oboes. The embellishing and emerging instrument itemsets heavily feature clarinets, as well as some additional oboes and flutes in the emerging instrument itemsets. [oboe 1]'s ubiquity as a dominating instrument and as an emerging instrument in this case study also causes some conviction values to drop to 0 and some leverage values to become negative. In comparison to the dominating instruments results in the flute case study, the most noticeable difference is that the support of [oboe 2, oboe 1] is lower than the support of [flute 2, flute 1]. This reflects what was seen in the previous chapter: when an oboe is dominating, it is less likely to be doubled by another oboe part than the flute was to be doubled by another flute.

The [clarinet 2, clarinet 1] itemset in the embellishing instruments once again has lift, leverage, and conviction scores that indicate a statistical dependence between the two instrument parts. The emerging instrument itemsets show some instances of statistical independence, with one negative correlation in [clarinet 1, oboe 1] due to oboe 1's ubiquity. The emerging instruments' items are similar to those found in the flute case study's emerging instrument itemsets. Both case studies include a flute, clarinet, and oboe itemset. All three instruments individually had high relative frequencies in the flute and oboe's case-study results in Chapter 2 (see sections 2.2.2.2 and 2.2.2.3). This ongoing presence demonstrates that the three instruments were grouped beyond the pairings demonstrated in the relative frequency analysis.

French horn

Table 3.5 shows the resulting itemsets for the French horn case study, as well as their accompanying association rule metrics. For the dominating, embellishing, and emerging instruments, the minimum support values were set to 0.40, 0.50, and 0.70, respectively.

Table 3.5: Association Rule Metrics for the Frequent Patterns in the French Horn Case Study

Blend Role	Itemset	Support	Ant.	Cons.	Support Ant.	Support Cons.	Conf.	Lift	Lev.	Conv.
Dominating Instruments	horn 1	0.932	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	horn 2, horn 1	0.858	horn 2	horn 1	0.872	0.932	0.984	1.056	0.045	4.358
Embellishing Instruments	clarinet 1, bassoon 2, bassoon 1	0.473	clarinet 1	bassoon 2, bassoon 1	0.601	0.709	0.787	1.109	0.046	1.361
	bassoon 2, bassoon 1	0.709	bassoon 2	bassoon 1	0.730	0.750	0.972	1.296	0.162	9.000
	clarinet 2, clarinet 1, bassoon 2, bassoon 1	0.459	clarinet 2	clarinet 1, bassoon 2, bassoon 1	0.588	0.473	0.782	1.652	0.181428	2.413
	bassoon 1	0.750	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Emerging Instruments	bassoon 1, horn 2, horn 1	0.840	bassoon 1	horn 2, horn 1	0.914	0.914	0.919	1.006	0.005	1.066
	horn 2, horn 1	0.914	horn 2	horn 1	0.939	0.975	0.974	0.998	−0.002	0.939
	bassoon 2, bassoon 1, horn 2, horn 1	0.822	bassoon 2	bassoon 1, horn 2, horn 1	0.871	0.840	0.944	1.123	0.090	2.831
	horn 1	0.975	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

The items found in the resulting frequent itemsets are the same ones that were found in the relative frequency analysis's French horn case study (see section 2.2.2.4). It should be noted that, whereas the previous analysis found that the French horn is often doubled by another French horn part, [horn 1] does not have a support value of 1 in any blend role. This means that there are some instances where other French horn parts are dominating without the first part. However, the support is still high enough to generate a conviction of nearly 1 and a near-zero leverage value for [horn 2, horn 1] in the emerging instrument blend role. In the dominating instruments, however, the confidence and conviction values of [horn 2, horn 1] are high. Although the lift and leverage show statistical independence, horn 2 has a high likelihood of appearing with horn 1.

Based on instances of lift, leverage, and conviction values which favour statistical dependence, there are a few itemsets which possess slight statistical dependence. The first is [bassoon 2, bassoon 1] in the embellishing instruments. Although its lift is close to 1, its leverage is positive and considerably higher than 0, especially when compared to other itemsets. The relationship between the two items is further cemented by the conviction value of 9. The [clarinet 2, clarinet 1, bassoon 2, bassoon 1] itemset in the embellishing instrument blend role also has a high leverage value and one of the more favourable lift values. This is an instance of two different instruments being doubled, which has not been seen in the previous itemsets. This phenomenon reappears in the emerging instruments's itemset [bassoon 2, bassoon 1, horn 2, horn 1], but has slightly different results due to horn 1's large presence in the emerging instruments.

The rest of the itemsets have two different instruments in their antecedent and consequent, including the emerging instrument pattern [bassoon 1, horn 2, horn 1]. The additional horn parts do not show a strong statistical dependence on the lone bassoon part. This is reflected in the near-zero leverage value and the lift and conviction values of approximately 1.

3.4 Discussion

Section 3.2.1 explains the process of selecting the minimum support value, and hints at the potential for different support thresholds depending on the blend role or case-study instrument. When introducing Table 3.1, it was noted that the minimum support for the dominating instruments was set to 0.20 to obtain more than one result. Similarly, the embellishing instruments needed a minimum support of 0.30 to produce more than one result. This is a significant difference when compared to the emerging instruments' minimum support of 0.60, which easily produced multiple results. There are two explanations for this difference. The first is applicable to the dominating instrument blend role. As explored in the previous chapter, dominating instrument groups rarely mix multiple instrument families together. This phenomenon was also seen in the dominating instrument itemsets in this analysis, which usually consisted of only one instrument part or two parts of the same instrument. This greatly limits the instruments that can be featured together in a dominating instrument grouping. For smaller instrument families, like the strings, this could have a large impact on the support values. This leads to the second explanation, which is applicable to all three blend roles. The previous chapter reported the average group size of each blend role. Dominating instrument groupings consist of, on average, two instrument parts. This is in comparison to the average of seven instrument parts for embellishing instrument groupings and the average of 13 instrument parts for emerging instrument groupings. These different averages result in different average transaction lengths, or number of items, for each grouping of instruments. With the size and family constraint on the dominating instruments, it is much more challenging to find itemsets in this blend role with high support values than for a blend role with lengthy transactions, like emerging instruments. These two issues become less prevalent in the case-study instruments, as the dominating and embellishing instrument association rule metrics benefit from the greater focus.

The featured instruments found in the larger blend roles align with the relative frequency findings in the previous chapter (see section 2.2.1). The most frequently used sections are violin 1 and violin 2 for dominating instruments, clarinet 1, clarinet 2, and bassoon 1 for embellishing instruments, and bassoon 1 and clarinet 1 for emerging instruments. These results also highlight the similarities between embellishing instruments in timbral augmentation blends and instruments used in timbral emergence blends. Both blend roles have a high average number of instrument parts, especially when compared to the dominating instruments (see section 2.3). As such, for each frequent itemset, it is likely that the antecedent's presence or absence from an annotation hardly has an impact on the consequent's presence. The embellishing or emerging instrument groups are large enough to include all instrument parts in the itemset. So, the support and confidence values for these itemsets are likely to be similar, which results in similar lift, leverage, and conviction values. This reasoning also explains why these values often indicate statistical independence (see section 3.1). The differences arise specifically in the exact support and confidence values. Given that timbral emergences most often occur when the orchestration plays *in tutti*, their frequent itemsets appear much more frequently than those in timbral augmentation blends, and thus have higher support values. The formula for confidence depends solely on support values, as explained in this chapter's background section (3.1), so it is impacted in turn.

The case-study results share many similarities with their relative frequency results. The most frequently occurring dominating, embellishing, and emerging instruments for each subject instrument are in line with those found in the previous chapter (see section 2.2.1). However, as opposed to the relative frequency of occurrence analysis, frequent pattern mining allowed for frequent itemsets beyond pairs, with some itemsets extending to three different instruments or to four instrument parts. Unsurprisingly, this occurred most often in the emerging instruments. These larger itemsets, however, usually feature two different instruments in total. For the embellishing instruments, this highlights the importance of doubling woodwind instruments. One example of this is the [clarinet 2, clarinet 1, bassoon 2, bassoon 1] frequent itemset in the French horn case study. Piston (1955) and Read (1979) have both mentioned that woodwinds possess certain acoustic weak points in their range. This could explain why the clarinet and bassoon had to be doubled by more than one section each. The oboe, flute, and clarinet combination, as mentioned in Adler (1982)'s treatise, appears frequently throughout the emerging instruments.

The association rule metrics provided more insight to the case-study patterns. Support is another form of measuring the relative frequency of occurrence, but the confidence, lift, leverage, and conviction values gave a clearer picture of the implications of these frequent patterns. Most notably, lift, leverage, and conviction detected the statistical dependence of itemsets that consist solely of multiple parts of the same instrument. In comparison, itemsets consisting of instrument parts of different instruments were largely statistically independent of one another. This is an important point to note when conducting a corpus analysis. Although certain instruments may have a high likelihood of being used together in a similar blend role, as indicated by a high support or confidence value, the final instrumentation remains the orchestrator's choice. As Piston (1955) and Read (1979) both noted, doubling an instrument using multiple parts may be a necessity due to an instrument's weaker acoustics. However, the choice of arranging certain instruments together seems to have been left up to the orchestrator.

Another motivation for this analysis was to determine blend strength's role on the detected frequent instrumentation patterns. With the minimum blend strength set to 2 (see section 3.2.1), the resulting itemsets were not surprising, considering the results from the previous chapter. When adjusting the blend strength to values between 0 and 3, there were no differences in the generated itemsets, apart from slight changes in their association metric values. The only major change occurred when the minimum blend strength was set to 4 or 5, which often resulted in no frequent itemsets being generated. This may be due to the nature of frequent patterns and the distribution of blend strengths across the annotations. Table 2.1 in Chapter 2 shows that most blend annotations are at a strength of 3 or 4. The results from the Shapiro-Wilk normality test on these data indicated that the blend strengths follow a normal distribution as well. So, if an itemset has a high enough support value, there is a greater likelihood that many of its annotations have a higher blend strength. Even if the itemset is featured in some low-strength annotations, it is most likely balanced out by another blend with a strength of 4 or 5. Although the inclusion of this pruning variable can be of use in future applications of this implementation of the FP-Growth algorithm, it did not have the expected impact in the current study.

There were limits to how much these association rule metrics could reveal for the case-study instruments. These metrics are built with the expectation that no item will appear in every single transaction. However, due to the nature of the case studies, the first part of the subject instrument often ends up appearing in every annotation. This is the cause of most negative correlations found in the case-study results. However, some of the emerging instrument itemsets show that adding additional instruments as the antecedent rectifies this problem. This creates a challenge for the analysis, as one must contextualize the statistics with the constant presence of one instrument part. This problem does not affect the blend roles, as there is no instrument part that is found in every annotation. Fortunately, the current study benefits from having multiple analysis techniques to rely on when examining OrchARD and our instrument case studies. So, the frequent pattern mining results were able to be interpreted and contextualized using results from the previous relative frequency analysis.

The similarities between the relative frequency results and the frequent pattern mining results show that both techniques are complementary analysis tools for corpus analysis. Although blend strength did not have the influence that was anticipated, the association rule metrics proved to be an insightful addition to relative frequencies of occurrence. One thing to note is that there were very few cases of strong statistical dependence between the different instrument sections, as shown by the consistently low leverage and conviction values. This may, however, be the result of limited data within the corpus or part of the musical reality. Such queries will be further addressed in Chapter 5 (see section 5.2).

Chapter 4

Experiment 3: Network Analysis of Timbral Augmentation Blends

Network science as a field has experienced significant growth in the 21st century so far. The rise of new technology, as well as the interdisciplinary nature of network science, meant that tools were created to analyze complex behaviour structures using network maps (Barabási, 2016). As a discipline, network science has seen applications in numerous fields, ranging from immunology to political science. Networks are of great interest due to the information provided by their structure and the ability to analyze complex relationships (Hevey, 2018). Barabási (2016) explained that networks provide different data with a common structure. For example, a network of Twitter users can be analyzed, and its structure can be compared to a network of club members, as they both fall under social networks. Networks have also been used in music research. Musicology has used it to explore social networks between composers and creative collaborators (Gleiser & Danon, 2003; Park et al., 2014; Park et al., 2015; Uzzi & Spiro, 2005). In music information retrieval (MIR), it has been used to detect listening patterns and improve music recommendation engines (Cano et al., 2006). As Hevey (2018) mentioned, networks have not been as pervasive in psychology, let alone in music psychology. This makes network science an interesting tool to use on OrchARD's annotations.

The aim of this chapter is to explore the interactions between different instruments in the timbral augmentation blend roles, and to explore potential subgroupings in the dominating and embellishing instruments. For instance, do particular dominating or embellishing instrument parts have certain roles within their timbral augmentation blends? First, background will be provided for network science concepts and algorithms relevant to this analysis (see section 4.1). After introducing these tools, the methodology (see section 4.2) and the results (see section 4.3) of a network analysis on OrchARD's timbral augmentation data will be addressed.

4.1 Background on network science

4.1.1 Terminology

A *network* is a graph structure consisting of points, known as *nodes* or *vertices*, that are connected by *edges*, or *links*. These links can be *directed*, meaning they point from one node to the other, or *undirected* (Newman, 2018). There are many different types of network structures. Directed networks, or *DiGraphs*, consist exclusively of directed edges. The edge's direction represents the dynamic of that one-way relationship. For instance, a directed link in a phone call network can represent node A calling node B. Figure 4.1 compares an undirected network with a DiGraph. A *weighted network* is one in which all edges have a weight value. An edge's weight represents some numerical value attached to that relationship. In the example of the phone call network, an edge's weight can represent the duration of a phone call between nodes A and B. Unweighted networks usually set all edge values to 1. A *bipartite network* divides the nodes into two disjoint sets, A and B, such that all edges connect the nodes of set A to the nodes of set B (Barabási, 2016). For example, a bipartite network can be used to represent a group of people at a restaurant, with edges connecting them to what food they had ordered. Figure 4.2 shows an example of a bipartite network. All of the network structures mentioned above will be used in this analysis.

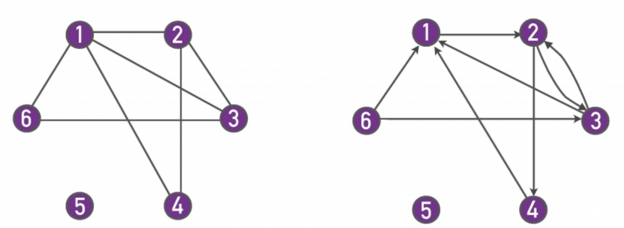

Figure 4.1 An example of an undirected network (left) and a directed network (DiGraph) (right). Taken from Barabási (2016)

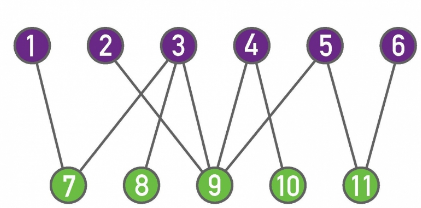

Figure 4.2 An example of a bipartite network. Taken from Barabási (2016)

Networks can also be represented mathematically using an *adjacency matrix*. The adjacency matrix is an n by n matrix, where n is the total number of nodes in the network. Each column indicates the nodes j that are connected to row i. Each entry A_{ij} is defined as follows:

$$A_{ij} = \begin{cases} 0, & \text{if there is no edge between } i \text{ and } j \\ w_{ij}, & \text{otherwise} \end{cases}$$

where w_{ij} is the weight of the edge between node i and node j (Newman, 2018). Adjacency matrices are a useful tool for determining the *spectra* of a network, which is a list of the adjacency matrix's *eigenvalues* in decreasing order. The spectra is often used to determine elements of the graph's structure, such as its connectedness, and the largest eigenvalue, also known as the *leading eigenvalue*, indicates the upper limit of the data's variance (Newman, 2018). Although adjacency matrices are not directly used in this chapter's network analysis, their eigenvalues play a role in the centrality metrics, or measures of node importance.

There are some basic network statistics that can be measured to gain further insight into a network's structure. *Degree* is one of the most commonly studied characteristics of graph structures. It is the sum of the edge weights connected to a node. In an unweighted network, all edge weights are equal to 1, so the degree is an indicator for a node's number of *neighbours*, or the nodes that a node is directly connected to. Nodes in a DiGraph can also have *in-degrees*, meaning the sum of the edges which point to the node in question, and *out-degrees*, the sum of the edges pointing away from the node (Barabási, 2016). Degrees can provide insight into the characteristics of a network's structure. For instance, do all nodes have approximately the same degree, or do most nodes have small degrees with a few high-degree nodes anchoring the network? Degrees can also be used to detect some important nodes, which will be discussed further when covering the centrality algorithms (see section 4.1.2).

A *path* is a set of edges that connect node A to node B. The shortest path of a network is important for certain network effects, like the *Small World effect*, in which two nodes in the largest networks have a rather short path separating them. A collection of nodes in a graph that have at least one path that connects them to any other node in the collection is called a *component*. A network that consists of only one component is said to be *connected* (Newman, 2018). Paths and edges, in general, are useful when looking for underlying group structures within a network. This topic will be approached in more detail when discussing community detection (see section 4.1.3).

4.1.2 Centrality

Large-scale networks with millions of nodes are of great interest to researchers. However, it can be daunting and inefficient to examine each node one at a time. Determining the most important nodes of the graph can guide researchers to key regions of the network and give greater insight into how the different elements of the data interact with one another. For this reason, a large portion of network science research has been dedicated to this task, known as *centrality.* Newman (2018) describes centrality as the answer to the following question: what nodes are the most important in this network? This question, though simple, can produce many different answers depending on one's definition of importance. As such, numerous algorithms have been proposed to address this question. This analysis uses two popular methods: *degree centrality* and *PageRank centrality.* Figure 4.3, taken from Needham and Hodler (2019), highlights the difference between degree centrality and PageRank centrality.

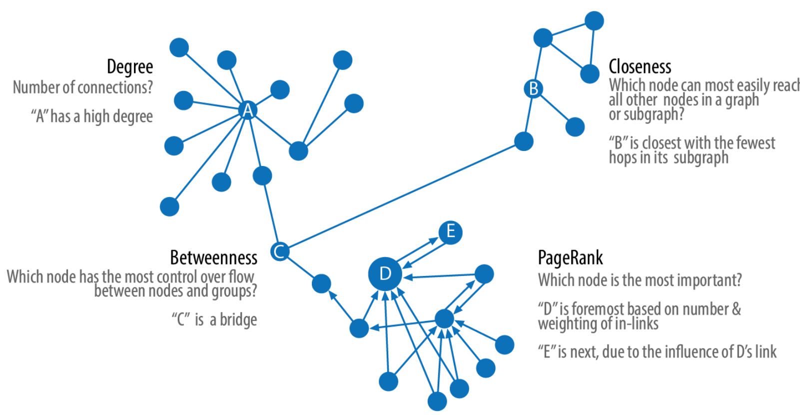

Figure 4.3 The difference between the different centrality algorithms. Assume that all edges have a weight of 1. Taken from Needham and Hodler (2019)

Degree centrality is the simplest and most intuitive of the centrality methodologies. In this case, the most important node is the node with the highest degree. A node's degree centrality, x_k, can be calculated as follows:

$$x_k = \sum_{i=1}^{n} a_{ik}$$

where n is the number of neighbours and a_{ik} is the weight of node k's neighbour i (Gao et al., 2014). Degree centrality depends on the node's degree value, meaning that every node's degree centrality value is greater than or equal to 0.

However, degree centrality's definition of importance raises some issues. In some network structures, having a high degree does not necessarily signify importance. It does not take into consideration what kind of neighbours a given node has. Is it attached to other high-degree nodes, or to nodes with no other neighbours? As Newman (2018) explains, importance should be about *who* you know, not just how many people you know.

PageRank is one of many solutions proposed to address degree centrality's shortcomings. Created for Internet search engines, most notably Google, it is a method that was initially used to rank every web page's importance. Web crawlers could then access and apply PageRank to all webpages without it being too computationally intense (Page et al., 1999). The equation is as follows:

$$x_k = \alpha \sum_j A_{ij} \frac{x_j}{d_j}$$

The PageRank of node k is a summation over the node's neighbours. The PageRank values for each node are initialized to the same value at the start of the calculation. For neighbour j, the weight of the link i-j is taken from the adjacency matrix A. It is then multiplied by the PageRank of each neighbour (x_j), then divided by neighbour j's out-degree, d_j. If d_j equals zero, then it is artificially set to 1 to avoid any divisions by zero. Once this multiplication is conducted for every neighbour j, these values are summed and then multiplied by *alpha*. *alpha* is a positive damping constant, and its value is influenced by the leading eigenvalue of the adjacency matrix. It is often to $\alpha = 0.85$ (Newman, 2018). The output is a probability distribution, meaning that the value is on a range from 0 to 1. The PageRank calculation depends on the PageRank value of all neighbours, meaning if one node's PageRank changes, the PageRank of each neighbouring node is updated.

These two centrality algorithms were carefully chosen to fit the research objectives of this chapter's network analysis. It was essential to select contrasting algorithms, as each one has its advantages and drawbacks. These two provide two contrasting definitions of what importance means: how many connections does a node have *and* what kind of connections does a node have. The common nodes in both sets of results, as well as the differences, will reveal further information about the nodes' functions and roles within their network.

4.1.3 Community detection

When analyzing larger data structures, like networks, clusters can be used to identify substructures and common functions among the different elements. These groupings, known as *communities* in network science, are groups of nodes that are more likely to be connected among themselves than to nodes in other communities. Barabási (2016) defines community strength using three types of communities:

i. *Cliques*, a subgraph in which all nodes are connected to each other.

ii. *Strong Communities*, a subgraph in which the nodes have more links to nodes within its community than with nodes in other communities.

iii. *Weak Communities*, in which the total internal degree of the subgraph is less than its external degree.

The quality of these communities can be evaluated using their modularity scores. A high modularity score, that is a positive score closer to 1, indicates an optimal partition. A positive modularity score that is closer to 0 indicates suboptimal partitions. A modularity score of 0 indicates an absence of communities, whereas a negative score indicates that every node is in its own community (Barabási, 2016). Modularity is thus a useful tool when evaluating the quality of partitions.

Network science researchers have proposed numerous strategies to detect communities when they are not predefined. Two prominent algorithms have been created with the thought of maximizing community modularity scores: the Clauset-Newman-Moore (CNM) algorithm and the Louvain algorithm. The two communities' methods will be explained using the simplified network in Figure 4.4.

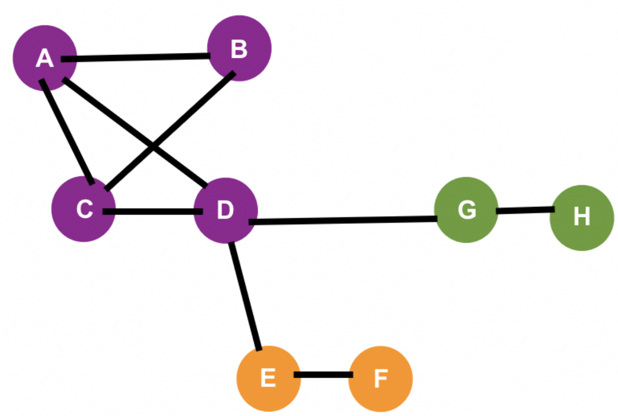

Figure 4.4 A simple undirected network with three groups

The CNM algorithm works as follows. When the algorithm begins, each node is assigned its own community. On each iteration, a node is moved around to each community and the modularity gain is recorded. If there is no move that increases the modularity, then the node remains in its current community. Otherwise, the node is added to the community that results in the highest modularity gain. For instance, suppose that at the beginning the CNM algorithm, each node in Figure 4.4 is in its own community. The algorithm begins with node A, and searches for a new community to add it to. Adding node A to the node F community would result in a small modularity gain. However, adding node A to the node C community would result in a much larger modularity gain, as they have more neighbours in common. Assuming that this produces the largest modularity gain of any potential move, node A is added to the node C community. This process continues until there is either no move produces a modularity gain or the highest gain falls below a user-defined threshold (Clauset et al., 2004). In computer science, these algorithms are known as *greedy algorithms*. The algorithm tries to maximize a certain value, in this case modularity, at every move. Although it can be effective in certain scenarios, it risks sacrificing a more strategic placement later for a larger increase in the moment (Cormen et al., 2022).

The Louvain algorithm, albeit similar to the CNM algorithm, takes a slightly different approach to address the weaknesses of greedy algorithms. After running the same procedure as the CNM algorithm, the Louvain algorithm then creates a new network in which each node is a community from the previous network. For example, assume that the colours in Figure 4.4 represent the communities detected by the CNM algorithm. The Louvain algorithm creates a new network of three nodes, with each one representing a different network. The CNM algorithm is then reapplied to this network. The Louvain algorithm repeats this process until the modularity gain falls below the defined threshold (Traag et al., 2019). This takes the greedy approach a step further and allows for some previous moves to be corrected as communities combine.

4.2 Method for network construction and analysis

In contrast with the analysis techniques presented in Chapters 2 and 3, this network analysis focuses heavily on the relationships and interactions between the instrument parts. The goal for this analysis technique is to explore some of the more intricate relationships within the dominating and embellishing instrument roles, as well as identifying potential subgroups within these blend roles. The visual representations and the different algorithms make networks an ideal analytical tool for OrchARD's data. In addition, the concept of analyzing the instrument parts in this manner makes sense when considering how orchestral instruments interact with one another. Composers and orchestrators are taught to view the ensemble as one entity, with the different musicians, as well as their instruments' acoustic and timbral properties, interacting with one another and moving as one (Piston, 1955). This interactive perspective of orchestras encourages an analysis using a web or network-like structure.

This chapter strictly focuses on timbral augmentation blends, which involve two blend roles: dominating instruments and embellishing instruments (see section 1.1.1). The previous two chapters also incorporated timbral emergence annotations to create a contrast between the different blend roles. Now that these differences have been established, the dynamic between dominating and embellishing instruments can be further investigated on its own. Previous chapters have clarified how certain instruments are assigned certain blend roles in timbral augmentation blends. Dominating instrument groups largely consist of multiple parts of the same instrument or of instruments with similar timbres. However, the instrumentation of embellishing instrument groups is still unclear, as frequent pattern mining generated roughly the same instruments for each case study. Before beginning this analysis, it was hypothesized that the subdivisions produced by community detection would categorize instrument nodes based on their subfunction within the blend role, specifically within the embellishing instruments blend role.

4.2.1 Network design

Although networks are not new to music research, orchestra networks and blend role networks are not as commonly studied. Some ideas for the design and procedure of this analysis were taken from previous work, particularly from Park et al. (2015) which created a network of Western Classical composers. Park et al.'s networks are social networks, meaning that they represent the interactions between individuals. In their case, those individuals are composers. Of the four broad network types outlined by Newman (2018) (technological, information, social, biological), orchestra networks and blend role networks were thought to be closest to social networks due to the interaction required with the musicians playing those instruments. For this reason, Park et al.'s design method for their networks was used as the basis for this chapter's network construction.

Eleven networks were created. In each network, a node represented an instrument part that was featured in whatever data subset was represented in the graph. A link was added between two nodes that appeared in at least one annotation together. The first ten networks focused on the two blend roles in timbral augmentation blends: dominating instruments and embellishing instruments. In these networks, two instrument parts that share an edge must have appeared in the same blend role together. Two networks (1, 2) encompassed all instrument parts featured either as dominating or embellishing instruments. The remaining eight (4–11) focused on the blend roles within the context of our four case-study instruments: violin, flute, oboe, and French horn. These networks were all weighted and undirected. The other network (3) was the Timbral Augmentation Network, which encompassed all instrument parts featured in the timbral augmentation annotations. To distinguish between the dominating and embellishing instruments, this graph was built as a weighted DiGraph. Dominating instrument nodes pointed to their embellishing instrument nodes. This means that, for the Timbral Augmentation Network, edges only existed between instruments in separate blend roles.

For all eleven graphs, the edges were weighted based on the average blend strength across all entries that featured two connected nodes. These weights were then normalized based on how frequently the combination appeared in the data. It did not make sense to attribute the same weight to a combination that is consistently strong and frequently seen and a combination that appeared only in one or two strong blends within the data. Chapter 3 did show that blend strength had a minimal effect on the frequency of occurrence of pairings (see section 3.4). However, when isolating each parameter, there were some instances in which the final results were stronger with only blend strength determining edge weight rather than frequency of occurrence. This was particularly the case for modularity values in the case-study instrument networks. Although the frequency of occurrence did ultimately have greater weight and influence on the final results, the strength of the blend formed by a pair of instrument nodes was included to strengthen the results.

The ten undirected graphs were constructed using bipartite projection. The technique was adapted from Park et al. (2015). The networks were initially built as bipartite graphs. One set of nodes consisted of all timbral augmentation annotations needed for the constraints of the graph. These nodes had edges which point to a second set of nodes, consisting of all instrument parts featured within the data or the case-study subset. Edges were weighted based on the blend strength associated with the annotation. From this, an undirected network was built from the set of instrument part nodes. Only instrument nodes with at least one edge pointing to it were featured, and undirected edges were formed to represent two instrument nodes that appear in at least one annotation together. In this method, the average weights were created during the bipartite projection process. In comparison, the Timbral Augmentation DiGraph was created manually from the data.

4.2.2 Procedure for network construction and analysis

The network creation and analysis were conducted in Python, primarily using features from the *NetworkX* library. Once the networks were created, their basic statistics were extracted. This was done to give the overall shape of the network without the use of visualizations. The information extracted at this stage included the number of nodes and edges, the number of components, the largest component size, the average degree, the node with the largest degree, and the heaviest edge. Next, the two centrality algorithms, degree centrality and PageRank centrality, were applied to each network. The top five most important nodes, as determined by each centrality measurement method, were recorded to provide a comparison with the network's node with the highest degree. Lastly, the two community detection algorithms, the CNM algorithm and the Louvain algorithm, were applied to the networks. This was done to detect potential subcategories within each blend role. The modularity calculator used to evaluate the quality of the community divisions was adjusted to favour larger partitions over smaller ones.

Visual representations were also created to take advantage of the networks' unique representation medium. These images also made it easier to view the communities. From Python, the networks were exported to `.gml` files and imported into *Gephi*,[1] an application that creates networks from `.gml` files. Within *Gephi*, the networks were spatially arranged using the *ForceAtlas* layout function. This function gives the nodes their own gravity. Nodes that are highly connected to each other attract each other and repel the less connected nodes. It should be noted that *ForceAtlas* does not take into account the community grouping that has been applied to a network. However, depending on the function used, there may be some overlap between the community groupings and the structure formed by *ForceAtlas*. Each instrument node was also colour-coded based on its CNM or Louvain community assignment to aid in the visualization of community groups. Edges are colour-coded as well. Links within the same community had the same colour as the community. However, for links with nodes in different communities, the link colour is a mixture of the two community colours. In practice, this results in many grey or brown links. If the community detection algorithms produced different partitions, then two visualizations were created to show each set of results.[2]

4.3 Network analysis results

Table 4.1 details the network statistics for the eleven networks. There are some differences that can be identified from this first glance at the networks. First, there is only one network that is disconnected, in the sense that there are some isolated nodes and multiple components: the Dominating Instruments Network. The larger of the two components includes 44 of the 46 instrument nodes, meaning that there are two instrument parts that are found in the same annotation but do not dominate with any other instrument part. The two remaining nodes, as will be discussed further in section 4.3.1.1, are harp 1 and harp 2. All other networks are fully connected. It should be noted, however, that this does not mean that every possible instrument part combination exists in the network's annotations.

[1] https://gephi.org/

[2] The Jupyter notebook and Gephi files are available upon request from stephen.mcadams@mcgill.ca.

Table 4.1: Network Statistics

Network Name	# of Nodes	# of Edges	# of Components	Largest Component Size	Average Degree	Largest Degree	Heaviest Edge
1. Dominating Instruments Network	46	227	2	44	4.084	Violin 2 (17.22)	Violin 2 - Violin 1
2. Embellishing Instruments Network	62	1202	1	62	37.466	Clarinet 1 (154.62)	Clarinet 2 - Clarinet 1
3. Timbral Augmentation Network	64	1366	1	64	16.212	Violin 2 (69.81)	Violin 1 - Bassoon 1
4. Violin Dominating Instruments Network	8	17	1	8	9.130	Violin 2 (17.22)	Violin 2 - Violin 1
5. Violin Embellishing Instruments Network	52	841	1	52	19.365	Oboe 1 (68.85)	Clarinet 1 - Clarinet 2
6. Flute Dominating Instruments Network	19	86	1	19	0.852	Flute 1 (5.05)	Flute 1 - Flute 2
7. Flute Embellishing Instruments Network	42	447	1	42	4.992	Clarinet 1 (20.22)	Clarinet 1 - Clarinet 2
8. Oboe Dominating Instruments Network	15	77	1	15	0.771	Oboe 1 (2.86)	Oboe 1 - Oboe 2
9. Oboe Embellishing Instruments Network	39	323	1	39	2.081	Clarinet 1 (8.36)	Clarinet 1 - Clarinet 2
10. French Horn Dominating Instruments Network	29	180	1	29	2.532	Horn 1 (12.36)	Horn 1 - Horn 2
11. French Horn Embellishing Instruments Network	47	559	1	47	6.798	Bassoon 1 (25.97)	Bassoon 2 - Bassoon 1

The largest degree and heaviest edge values have considerable overlap, as the node with the highest degree is often included in the heaviest edge. The two exceptions to this phenomenon are the Timbral Augmentation Network and the Violin Embellishing Instruments Network. The former lists violin 2 as the node with the highest degree, whereas the violin 1-bassoon 2 edge has the heaviest weight. This is a testament to violin 2's versatility in comparison to violin 1. However, versatility does not necessarily translate to high strength. This will be explored further when examining each network's centrality results.

4.3.1 Analysis of blend roles

Dominating Instruments Network

Figure 4.5 shows the visualization of the Dominating Instruments Network. The network visualization is colour-coded to show the different communities defined by the community detection algorithms. In this case, the CNM and Louvain algorithms produced the same partitions, so only one figure is shown. The thickness of the links is proportional to the edge's weight. A thicker edge represents a heavier edge. Table 4.2 then displays the detailed results for the centrality and community detection algorithms.

Table 4.2: Centrality & Community Detection Results for Dominating Instruments Network

Centrality Results		
Rank	Degree Centrality	PageRank Centrality
1.	Horn 1 (0.600)	Horn 1 (0.062)
2.	Horn 2 (0.600)	Horn 2 (0.062)
3.	Trumpet 2 (0.511)	Flute 1 (0.052)
4.	Trumpet 1 (0.511)	Trumpet 2 (0.046)
5.	Horn 4 (0.422)	Flute 2 (0.046)
Community Results		
Group Number	Louvain Communities	CNM Communities
1. (pink)	Trombone 1, Horn 2, Horn 5, Trumpet 3, Alto Trombone, Trombone 2, Tuba 2, Cornet 2, Trumpet 4, Trumpet 2, Horn 1, Horn 6, Tuba 1, Horn 4, Horn 7, Cornet 1, Horn 3, Bass Trombone, Trombone 3, Trumpet 1, Euphonium 2	Trombone 1, Horn 2, Horn 5, Trumpet 3, Alto Trombone, Trombone 2, Tuba 2, Cornet 2, Trumpet 4, Trumpet 2, Horn 1, Horn 6, Tuba 1, Horn 4, Horn 7, Cornet 1, Horn 3, Bass Trombone, Trombone 3, Trumpet 1, Euphonium 2
2. (green)	English Horn, Clarinet 3, Clarinet 1, Piccolo Clarinet, Piccolo 1, Flute 3, Bassoon 3, Bass Clarinet, Oboe 1, Flute 2, Flute 1, Flute 4, Contrabassoon, Clarinet 2, Bassoon 2, Bassoon 1, Oboe 3, Oboe 2	English Horn, Clarinet 3, Clarinet 1, Piccolo Clarinet, Piccolo 1, Flute 3, Bassoon 3, Bass Clarinet, Oboe 1, Flute 2, Flute 1, Flute 4, Contrabassoon, Clarinet 2, Bassoon 2, Bassoon 1, Oboe 3, Oboe 2
3. (orange)	Violin 2, Viola, Cello, Violin 1, Bass	Violin 2, Viola, Cello, Violin 1, Bass
4. (blue)	Harp 2, Harp 1	Harp 2, Harp 1

Figure 4.5 shows the two components of the network, which were first noted in Table 4.1. The two-node component consists of harp 1 and harp 2. They are used sparingly as dominating instruments in OrchARD's data, only appearing 6 and 3 times, respectively. This is the only network, out of all eleven networks, that is disconnected. Some nodes are only connected to one or two nodes, like the bass clarinet or tuba 2. However, there is still a path that connects them to the other nodes in this network. They are not isolated from all other instruments in the same way as harp 1 and harp 2. The different communities in the larger component do have some edges that connect them. In comparison to some of the edges in the string instrument community, these

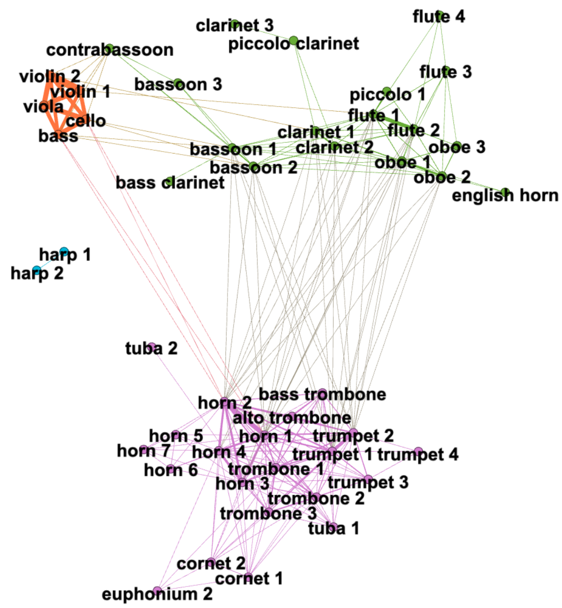

Figure 4.5 Dominating Instruments Network with Communities Formed by the CNM Algorithm and the Louvain Algorithm

inter-community edges are faint. Although there are some instances of instruments of different families forming dominating instrument groupings, it is a rare occurrence.

The degree centrality and PageRank centrality algorithms produced similar results. The top five nodes with the highest degree centrality were all brass instruments. More specifically, the most important nodes were different horn and trumpet parts. The PageRank centrality results also included two flute parts, which are some of the most commonly used dominating instruments in the woodwind family, as seen in Chapter 2. The centrality results show a stark difference from the highest degree and heaviest edge results in Table 4.1. Violin 2 has the highest degree, and violin 2-violin 1 has the heaviest edge. Edge weights depend on how frequently the two instrument parts appear, meaning that these two metrics reflect the frequency of occurrence. So, these results are in accordance with what was seen in past chapters (see sections 2.2.1.1 and 3.3.1.1). Despite their high frequency of use, the bowed string instruments do not have many links with non-string instrument nodes. This can be observed in the upper-left corner of Figure 4.5. The bowed strings have heavy edges among themselves but not with other instrument families. The horn 2 node, in comparison, has links with many different nodes both in the brass family and in other instrument families. It appears that the centrality metrics favour instruments with greater versatility as opposed to high blend strength and frequent use.

Both community detection algorithms produced the same four communities, achieving a modularity score of 0.787. This indicates strong communities and optimal partitioning. *ForceAtlas*'s partitions are also distinctly separate in both figures. The instrument nodes in this network are grouped according to their instrument family. Harp 2 and harp 1 are not included with the string instruments, probably because they are in a separate component. Previous chapters mentioned that dominating instrument groups tend to consist of instruments from the same family, and these community detection results confirm this further.

Embellishing Instruments Network

Table 4.3 shows the results of the centrality algorithm and the community detection algorithm for the Embellishing Instruments Network. Figures 4.6 and 4.7 show the *Gephi*-produced visualizations of the network, showing both the CNM and Louvain communities, respectively.

Table 4.3: Centrality & Community Detection Results for the Embellishing Instruments Network

Centrality Results		
Rank	Degree Centrality	PageRank Centrality
1.	Flute 1 (0.967)	Clarinet 1 (0.058)
2.	Clarinet 1 (0.967)	Clarinet 2 (0.058)
3.	Bassoon 1 (0.967)	Bassoon 1 (0.055)
4.	Clarinet 2 (0.967)	Oboe 1 (0.054)
5.	Oboe 1 (0.951)	Bassoon 2 (0.053)
Community Results		
Group Number	Louvain Communities	CNM Communities
1. (orange)	Clarinet 3, Horn 5, Piccolo Clarinet, Flute 3, Oboe 4, Bassoon 3, Trumpet 4, Horn 6, Flute 4, Horn 7, Piccolo 2, Triangle, Oboe 3	Clarinet 3, Horn 5, Piccolo Clarinet, Flute 3, Oboe 4, Bassoon 3, Horn 6, Flute 4, Horn 7, Piccolo 2, Triangle, Oboe 3
2. (blue)	Tubular Bells, Xylophone, Violin 2, Viola, Cello, Glockenspiel, Alto Saxophone 1, Gong, Harp 2, Violin 1, Bass, Harp 1	Tubular Bells, Xylophone, Violin 2, Viola, Cello, Glockenspiel, Alto Saxophone 1, Gong, Harp 2, Violin 1, Bass, Harp 1, Bass Drum, Bass Clarinet, Tam-Tam, Contrabassoon, Cymbals, Snare Drum
3. (pink)	Trombone 1, Horn 2, Trumpet 3, Trombone 2, Tuba 2, Cornet 2, Bass Drum, Bass Clarinet, Bass Horn 1, Trumpet 2, Horn 1, Timpani 1, Solo Violin, Tuba 1, Horn 4, Bass Horn 2, Tam-Tam, Contrabassoon, Cornet 1, Horn 3, Trombone 3, Cymbals, Trumpet 1	Trumpet 4, Trombone 1, Horn 2, Trumpet 3, Trombone 2, Tuba 2, Cornet 2, Trumpet 2, Horn 1, Timpani 1, Tuba 1, Horn 4, Cornet 1, Horn 3, Trombone 3, Trumpet 1, Alto Trombone, Bass Trombone
4. (green)	English Horn, Clarinet 1, Alto Trombone, Piccolo 1, Snare Drum, Oboe 1, Celesta, Flute 2, Flute 1, Clarinet 2, Bassoon 2, Bass Trombone, Bassoon 1, Oboe 2	Bass Horn 1, Solo Violin, Bass Horn 2, English Horn, Clarinet 1, Piccolo 1, Oboe 1, Celesta, Flute 2, Flute 1, Clarinet 2, Bassoon 2, Bassoon 1, Oboe 2

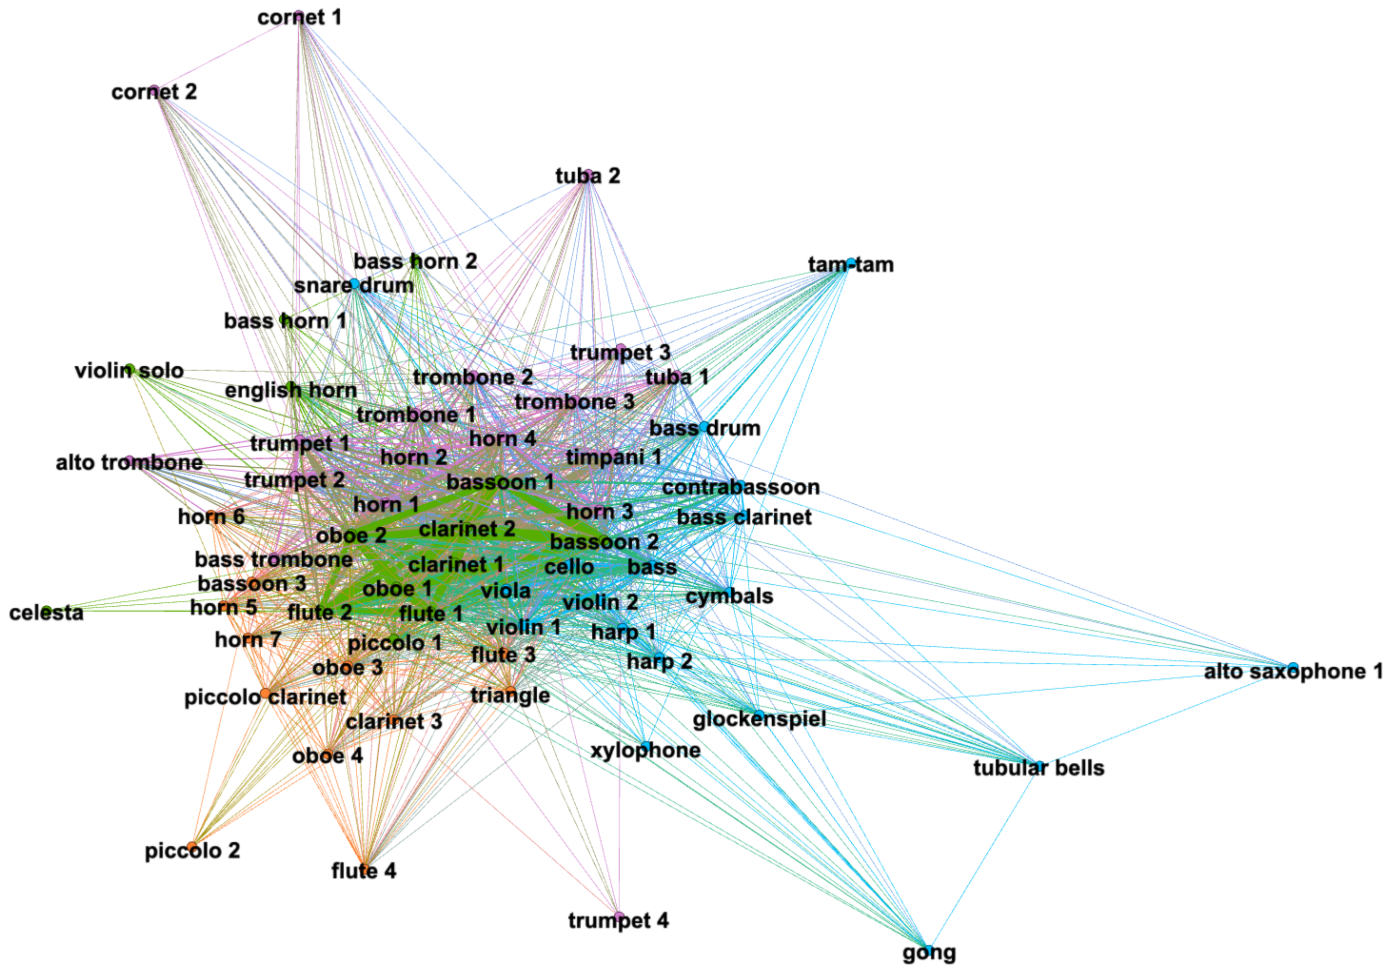

Figure 4.6 Embellishing Instruments Network with Communities Formed by the CNM Algorithm

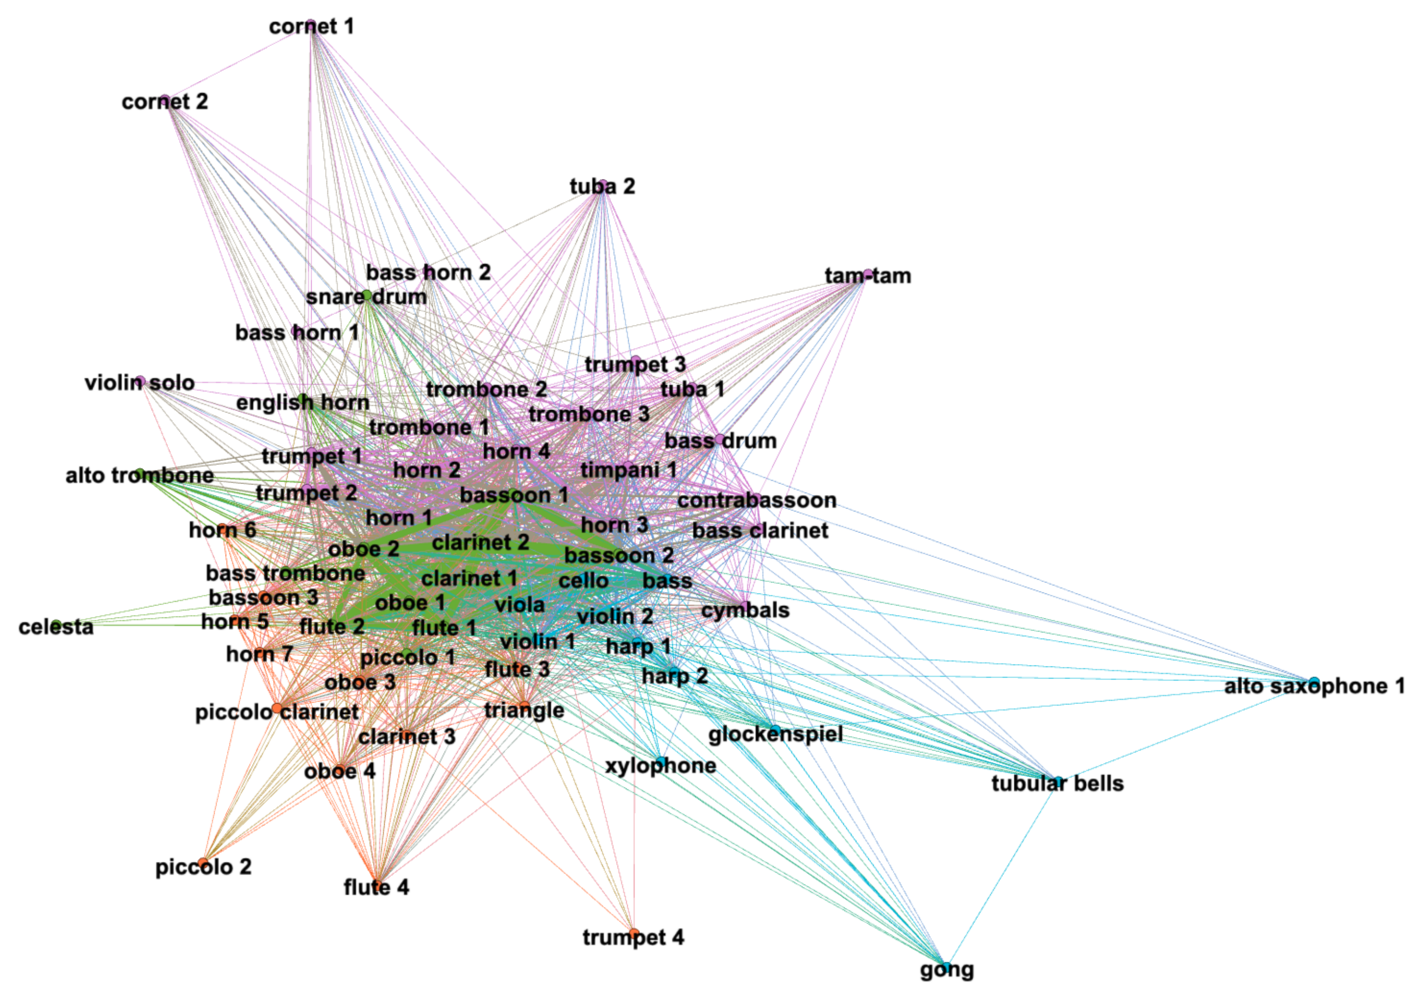

Figure 4.7 Embellishing Instruments Network with Communities Formed by the Louvain Algorithm

The Embellishing Instruments Network's spatial structure is vastly different from that of the Dominating Instrument Network. Rather than having distinctly separated communities, this network has a more web-like shape. Based on how *ForceAtlas* arranges nodes, the cluster of instruments in the center are all highly connected to each other and the nodes along the outer perimeter are less frequently used. Upon closer inspection, the nodes at the centre of the network have been previously reported to be frequently used embellishing instruments in the previous chapters. Instruments, such as the clarinet and the bassoon, have multiple nodes in the centre. The nodes around the edge are percussion instruments, like the tam-tam and the gong, or rarely used instruments in their respective family, like alto saxophone 1 or cornet 1.

The centrality results are slightly different between the two algorithms. Both methods revealed woodwind instruments to be the five most important nodes. However, degree centrality had a greater diversity in the woodwind instruments chosen than did PageRank. Flute 1 is the most important node according to degree centrality, whereas clarinet 1 is the most important instrument part according to PageRank. Based on the different formulae employed by the algorithms, it can be inferred that flute 1 has more neighbours than clarinet 1, but clarinet 1 has more important connections than flute 1. This shows that clarinet 1 is more versatile than flute 1. That is, clarinet 1 is paired with a greater variety of instrument parts in an embellishing instruments group than flute 1. Additionally, clarinet 1 also has the highest degree in the Embellishing Instruments Network. In this network's case, the most versatile instrument part is also frequently used in strong timbral augmentation blends, as the results from the centrality algorithms include the nodes with the highest degree and heaviest edge.

The community detection algorithms produced suboptimal partitions. The Louvain communities have a modularity of 0.254 and the CNM partitions have a modularity of 0.247. The two sets of communities have some slight differences. The placement of some less frequently used instruments, such as bass horn 1 or tam-tam, varies between the two algorithms. However, most of the instrument parts are placed in the same communities in both the Louvain and CNM algorithms. Group 1 consists of less frequently used wind instrument parts and the triangle. This group is located in the bottom left corner of Figures 4.6 and 4.7. Group 2 includes string instruments and less commonly used percussion instruments. It is in the bottom right corner of the figures. Group 3 consists mainly of brass instruments and timpani, and is found in the upper part of the network visualizations. The fourth group largely consists of different frequently used woodwind instruments, all found at the centre of the network. The communities appeared to be formed on the basis of how frequently the instrument part is used rather than distinct functions within the broader category of embellishing instruments. This could explain the low modularity score for both sets of partitions.

Timbral Augmentation Network

Table 4.4 presents the results for the centrality algorithms and the partitions for both community detection algorithms when applied to all of the instrument parts featured in timbral augmentation blends. The *Gephi* visualizations for this network are shown in Figures 4.8 and 4.9. Due to the large number of nodes and edges in this network, the arrows on the edges are minuscule. For this reason, a simplified version of this network without the community detection colour-scheme can be found in Figure 4.10.

Table 4.4: Centrality & Community Detection Results for the Timbral Augmentation Instruments Network

Centrality Results		
Rank	Degree Centrality	PageRank Centrality
1.	Violin 1 (1.476)	Bassoon 1 (0.070)
2.	Violin 2 (1.429)	Bassoon 2 (0.064)
3.	Viola (1.381)	Clarinet 1 (0.048)
4.	Cello (1.365)	Viola (0.047)
5.	Horn 2 (1.349)	Cello (0.047)
Community Results		
Group Number	Louvain Communities	CNM Communities
1. (light green)	Piccolo 1, Flute 1, Oboe 1, Oboe 2, Clarinet 1, Clarinet 2, Alto Saxophone 1, Horn 3, Trumpet 2, Trombone 1, Triangle, Glockenspiel, Xylophone, Solo Violin, Violin 2	Piccolo 1, Flute 1, Oboe 1, Oboe 2, Clarinet 1, Clarinet 2, Alto Saxophone 1, Horn 3, Trumpet 2, Trombone 1, Triangle, Glockenspiel, Xylophone, Solo Violin, Violin 2, Flute 3, Flute 4, Oboe 3, Oboe 4, Clarinet 3, Piccolo Clarinet, Horn 5, Horn 6, Horn 7, Trumpet 3, Trumpet 4, Horn 4, Piccolo 2, Bass Horn 1, Bass Horn 2
2. (orange)	Flute 3, Flute 4, Oboe 3, Oboe 4, Clarinet 3, Piccolo Clarinet, Horn 5, Horn 6, Horn 7, Trumpet 3, Trumpet 4	Horn 3, Trumpet 3, Bassoon 3, Contrabassoon, Trumpet 1, Cornet 1, Cornet 2, Trombone 2, Trombone 3, Euphonium 2, Tuba 1, Tuba 2, Tam-Tam, Gong, Cymbals, Bass Drum, Tubular Bells, Bass
3. (pink)	Bassoon 3, Contrabassoon, Horn 4, Trumpet 1, Cornet 1, Cornet 2, Trombone 2, Trombone 3, Euphonium 2, Tuba 1, Tuba 2, Tam-Tam, Gong, Cymbals, Snare Drum, Bass Drum, Tubular Bells	Snare Drum, Bass Clarinet, Bassoon 1, Horn 1, Bass Horn 1, Vibraphone, Celesta, Violin 1, Viola, Flute 2, English Horn, Bassoon 2, Horn 2, Alto Trombone, Bass Trombone, Timpani 1, Harp 1, Harp 2, Cello
4. (dark green)	Piccolo 2, Bass Clarinet, Bassoon 1, Horn 1, Bass Horn 1, Vibraphone, Celesta, Violin 1, Viola	N/A
5. (blue)	Flute 2, English Horn, Bassoon 2, Horn 2, Bass Horn 2, Alto Trombone, Bass Trombone, Timpani 1, Harp 1, Harp 2, Cello, Bass	N/A

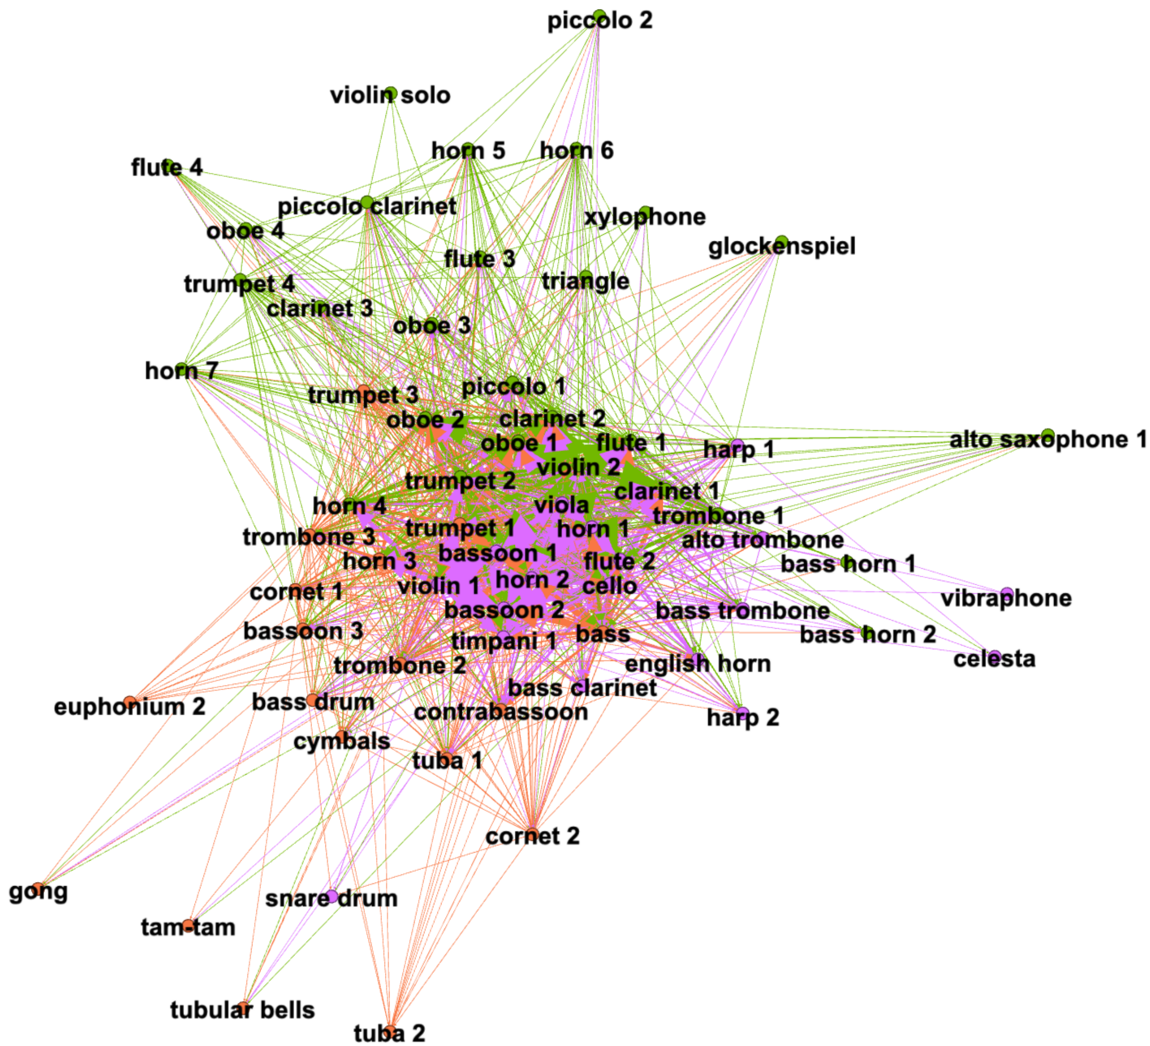

Figure 4.8 Timbral Augmentation Network with Communities Formed by the CNM Algorithm

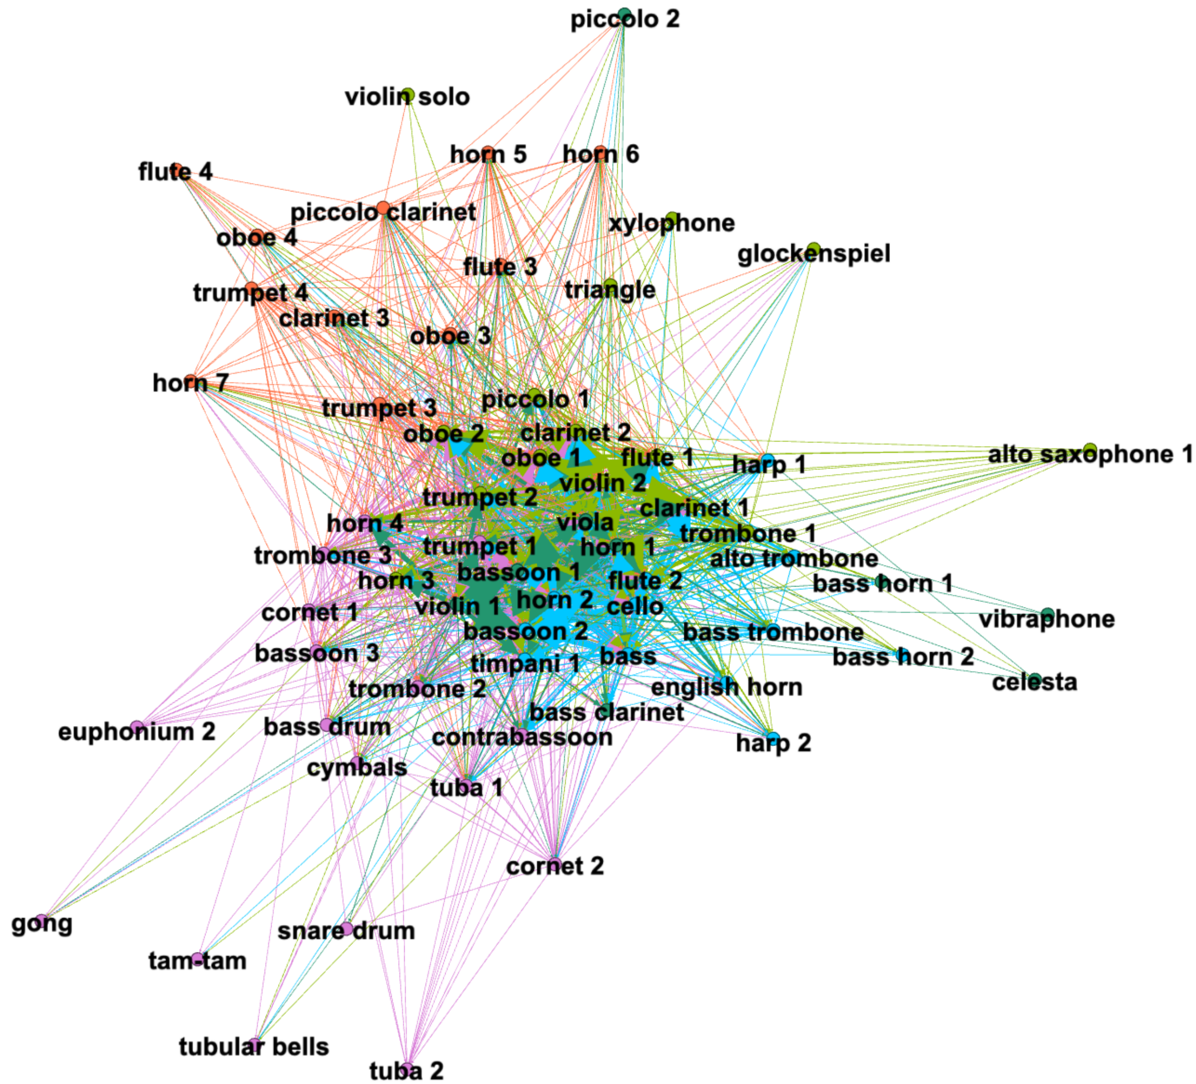

Figure 4.9 Timbral Augmentation Network with Communities Formed by the Louvain Algorithm

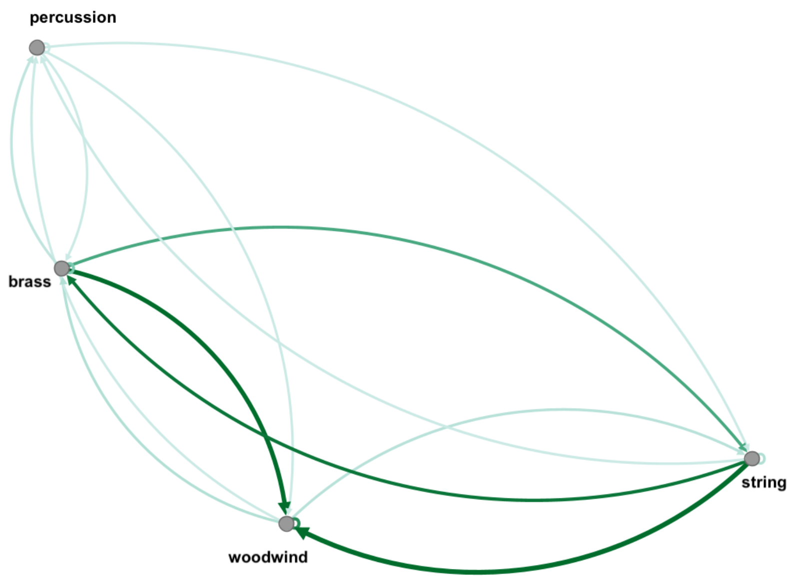

Figure 4.10 Simplified Timbral Augmentation Network using only Instrument Family Nodes

The Timbral Augmentation Network is the only DiGraph of the eleven networks produced for this analysis. As such, some of its results diverge from the results of other networks. Nevertheless, its structure created by *ForceAtlas* shows similar web structure to what was observed in the Embellishing Instruments Network. At its centre, there are different frequently used dominating and embellishing instrument parts. The outer edge consists of the less frequently used instrument nodes, regardless of blend role.

The simplified network in Figure 4.10 gives further insight into these relationships on an instrument family level. There are 16 edges, meaning that every instrument family has an edge pointing to and from every other instrument family node. Nevertheless, the edge weights in this network provide further insight. The gradient on an edge represents its weight, with darker edges being heavier. The thick edges pointing from the brass and the strings to the woodwinds depict the woodwinds' tendency to embellish in timbral augmentations. The edges going to and from the percussion node are incredibly pale, indicating the family's infrequency of use in timbral augmentation blends. The arrows pointing between the brass and string families are of particular interest. They are of similar weight, but the string-to-brass edge is slightly darker than the brass-to-string edge. This means that brass instruments embellish string instruments more often than string instruments embellish brass instruments.

The centrality algorithms produced vastly different results. The degree centrality's most important nodes consist of string instruments and the second French horn part. PageRank, on the other hand, deemed that bassoons, clarinet 1, viola, and cello were the most important instrument nodes. This means that the degree centrality nodes, which appear to favour instruments that are frequently used as dominating instruments, generally have more connections. On the other hand, the PageRank nodes, which consist of common embellishing instruments, have more important connections. This may be due to the construction of this network. Edges point from dominating instruments to embellishing instruments. This means that instrument nodes that are used mostly as embellishing instruments will have a higher in-degree, or a higher number of edges pointing towards them than away from them. PageRank may be identifying this as a metric for importance, which explains why it is only listing embellishing instruments as its top five most important nodes.

The communities created for the Timbral Augmentations Network are the weakest of the eleven networks. The Louvain algorithm and the CNM algorithm produced partitions with modularity scores of 0.178 and 0.236, respectively. This means that the communities produced are suboptimal. The two algorithms also did not produce the same number of communities. The CNM algorithm produced three partitions, whereas the Louvain algorithm produced five. As the low modularity scores hinted at, the communities formed seemingly have little in common with each other. As can be seen in Table 4.4 and the two figures (4.8 and 4.9), the groupings consist of instrument parts from different spatial regions of the network, from different instrument families, and from different frequencies of occurrence. For instance, group 1 in the Louvain communities consists of frequently occurring woodwinds, like flute 1 and oboe 1, and less frequently occurring instruments, like the xylophone and solo violin. The heterogeneity within the partitions could be due to the network's directed edges. These edges add a layer of complexity that is not present in the other networks, making it more challenging for these algorithms to produce optimal communities.

4.3.2 Analysis of case studies

Violin

The first network presented for this case study is the Violin Dominating Instruments Network. The centrality and community detection results are presented in Table 4.5. Figure 4.11 shows the visualization of this network.

Table 4.5: Centrality & Community Detection Results for the Violin Dominating Instruments Network

Centrality Results		
Rank	Degree Centrality	PageRank Centrality
1.	Violin 2 (0.857)	Violin 2 (0.226)
2.	Violin 1 (0.857)	Violin 1 (0.222)
3.	Viola (0.714)	Viola (0.183)
4.	Bass (0.714)	Cello (0.163)
5.	Cello (0.714)	Bass (0.147)
Community Results		
Group Number	Louvain Communities	CNM Communities
1.	Clarinet 1, Violin 2, Viola, Cello, Flute 1, Violin 1, Contrabassoon, Bass	Clarinet 1, Violin 2, Viola, Cello, Flute 1, Violin 1, Contrabassoon, Bass

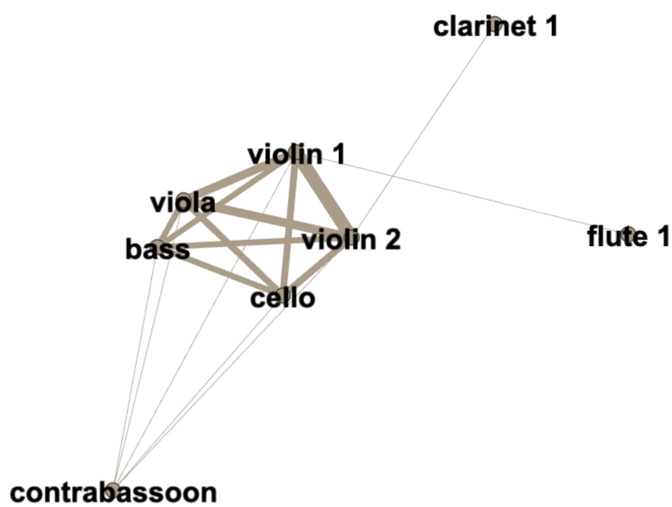

Figure 4.11 Violin Dominating Instruments Network with Communities Formed by the CNM Algorithm and the Louvain Algorithm

This is the smallest network, consisting of only eight nodes. This could explain why neither community detection algorithm was able to form partitions. Clarinet 1, flute 1, and contrabassoon only occur in one annotation and were omitted from the relative frequency analysis in Chapter 2 for that reason. Their placement within this network further justifies that decision. The contrabassoon shares edges with the five string instruments, but clarinet 1 and flute 1 are connected by a single edge to one of the string instruments. Unsurprisingly, the string instruments are clustered in the network's centre with heavy edges to connect them.

The centrality algorithms produce roughly the same results, ranking the string instruments as the most important instruments. It is important to note that violin 2 is the most important node in both centrality measurements and is also the node with the highest degree in this network. However, the frequent pattern mining analysis and the relative frequency analysis in the previous chapters determined that violin 1 occurs more frequently. Violin 1 appears on its own frequently, as shown in Figure 2.16 in Chapter 2. This means that, although it has a strong link with violin 2, it is not as strongly connected to the other string instrument nodes as the second violin part.

Table 4.6 shows the results of the centrality algorithms and the partitions created by the two community detection algorithms when applied to the Violin Embellishing Instruments Network. Figures 4.12 and 4.13 then show the network visualizations produced in *Gephi*.

Table 4.6: Centrality & Community Detection Results for the Violin Embellishing Instruments Network

Centrality Results		
Rank	Degree Centrality	PageRank Centrality
1.	Oboe 1 (1.000)	Oboe 1 (0.061)
2.	Flute 1 (1.000)	Flute 1 (0.060)
3.	Oboe 2 (0.980)	Clarinet 1 (0.059)
4.	Clarinet 1 (0.961)	Clarinet 2 (0.058)
5.	Flute 2 (0.961)	Oboe 2 (0.056)
Community Results		
Group Number	Louvain Communities	CNM Communities
1. (dark green)	Violin 2, Viola, Cello, Violin 1, Triangle, Bass	Violin 2, Viola, Cello, Violin 1, Triangle, Bass, Horn 2, Horn 1, Horn 4, Contrabassoon, Horn 3, Cymbals, Cornet 2, Trumpet 2, Timpani 1, Cornet 1, Trumpet 1
2. (pink)	English Horn, Trombone 2, Bass Drum, Bass Clarinet, Harp 2, Tuba 1, Harp 1, Trombone 3, Horn 2, Horn 1, Horn 4, Contrabassoon, Horn 3, Cymbals	English Horn, Trombone 2, Bass Drum, Bass Clarinet, Harp 2, Tuba 1, Harp 1, Trombone 3, Clarinet 3, Horn 5, Piccolo Clarinet, Flute 3, Oboe 4, Bassoon 3, Horn 6, Flute 4, Horn 7, Piccolo 2, Oboe 3
3. (light green)	Clarinet 1, Piccolo 1, Snare Drum, Oboe 1, Celesta, Flute 2, Flute 1, Clarinet 2, Bassoon 2, Bassoon 1, Oboe 2	Clarinet 1, Piccolo 1, Snare Drum, Oboe 1, Celesta, Flute 2, Flute 1, Clarinet 2, Bassoon 2, Bassoon 1, Oboe 2, Bass Horn 1, Bass Horn 2
4. (orange)	Cornet 2, Trumpet 2, Timpani 1, Cornet 1, Trumpet 1, Bass Horn 1, Bass Horn 2, Trombone 1, Alto Trombone, Bass Trombone	Trombone 1, Alto Trombone, Bass Trombone
5. (blue)	Clarinet 3, Horn 5, Piccolo Clarinet, Flute 3, Oboe 4, Bassoon 3, Horn 6, Flute 4, Horn 7, Piccolo 2, Oboe 3	N/A

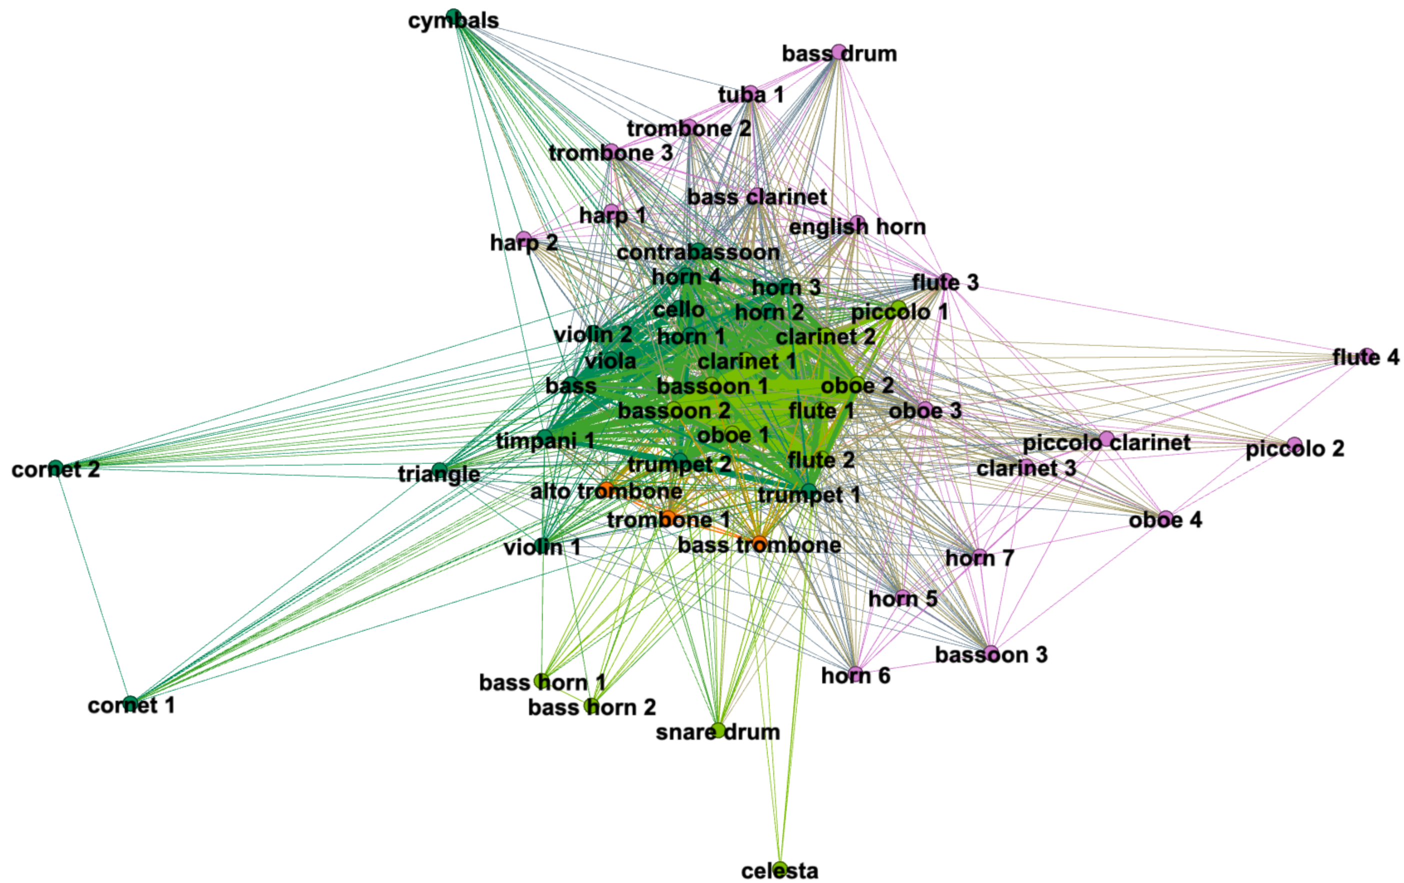

Figure 4.12 Violin Embellishing Instruments Network with Communities Formed by the CNM Algorithm

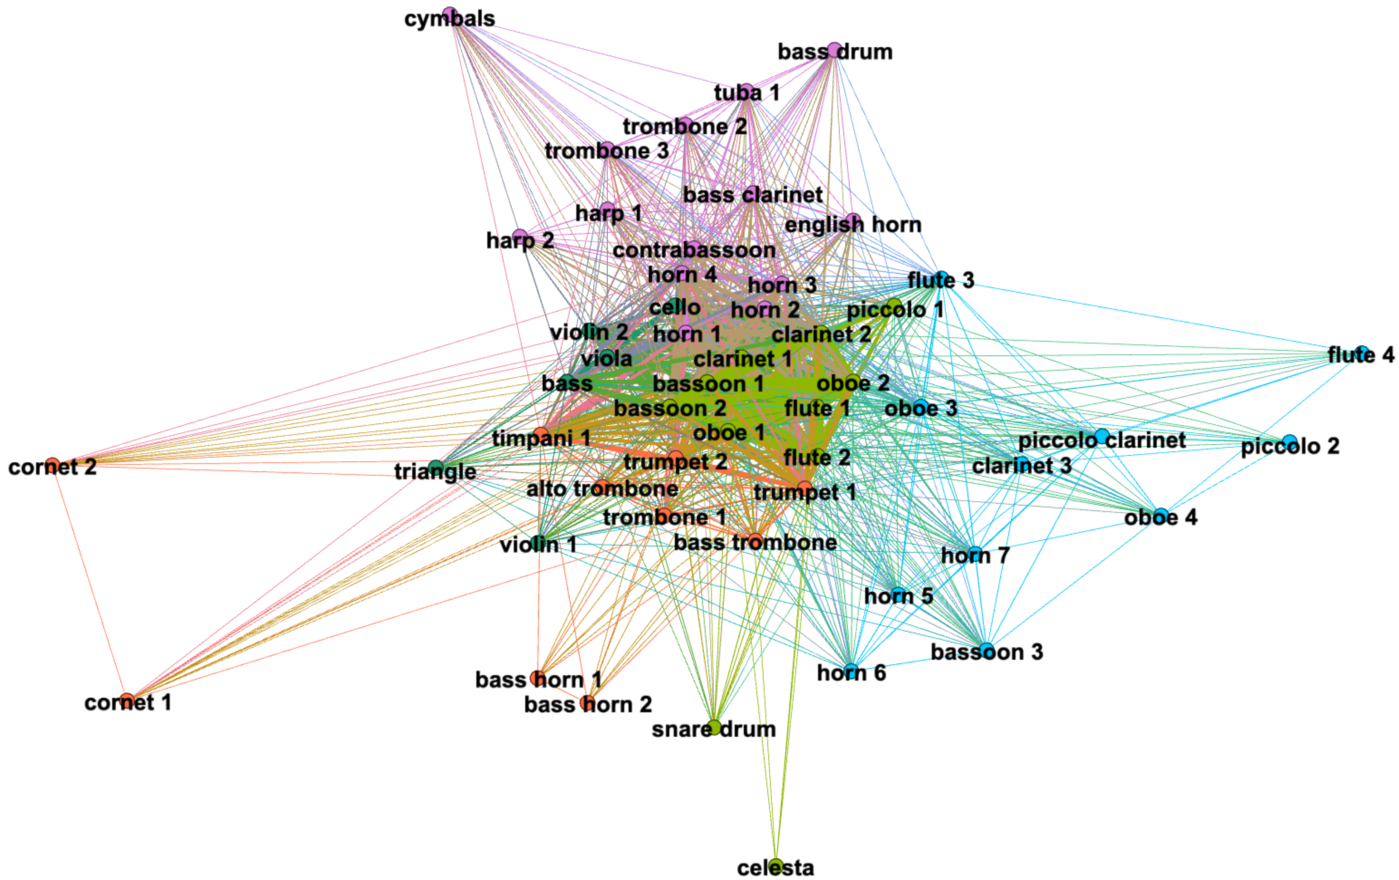

Figure 4.13 Violin Embellishing Instruments Network with Communities Formed by the Louvain Algorithm

Similarly to the Embellishing Instruments Network, the Violin Embellishing Instruments Network also has a web-like structure. Nodes with higher degrees, which were also deemed to be the most important nodes by both degree centrality and PageRank centrality, are clustered in the center. The less used instrument nodes are located along the edge of the network.

The Louvain algorithm and the CNM algorithm produced slightly different groups. Their partitions remain suboptimal, achieving modularity values of 0.226 and 0.260, respectively. The modularity function used was instructed to favour larger communities, which explains why CNM scored a higher modularity than Louvain. The first group indicated by Louvain largely consisted of string instruments, with the CNM's group also including brass instruments, as well as additional woodwind and percussion instruments. The second group produced by both algorithms has a less specific combination of instruments across all four instrument families. Group three consists of frequently used woodwind instruments, with some additional percussion. The CNM algorithm also included the bass horns in its third group. The fourth group in both algorithms consists solely of brass instruments. The additional fifth group produced by the Louvain algorithm consists of additional woodwind and brass instruments, which all occur less frequently than other instrument parts in their respective families.

There are some common traits within each community. The first and fourth groups seem to be partitioned on the basis of their instrument family. However, this division is influenced by additional noise, such as the presence of a triangle in group 1 or the additional brass instruments added to CNM's group 1. The other groups are partitioned similarly to the Embellishing Instruments Network, in that they are divided based on how frequently they occur. This explains why the brass and woodwinds are split across different groups. The algorithms also have difficulty determining how to place the percussion instruments, as they are so infrequently used.

Flute

Table 4.7 shows the centrality results and the partitions produced for the Flute Dominating Instruments Network. Figure 4.14 shows the network visualization, with the nodes colour-coded according to how each community detection algorithm created its groupings. As CNM and Louvain produced the same groupings, there is only one figure present.

Table 4.7: Centrality & Community Detection Results for the Flute Dominating Instruments Network

Centrality Results		
Rank	Degree Centrality	PageRank Centrality
1.	Flute 1 (1.000)	Flute 1 (0.262)
2.	Flute 2 (0.944)	Flute 2 (0.228)
3.	Oboe 1 (0.722)	Oboe 1 (0.070)
4.	Oboe 2 (0.722)	Oboe 2 (0.061)
5.	Horn 1 (0.611)	Clarinet 1 (0.049)
Community Results		
Group Number	Louvain Communities	CNM Communities
1. (orange)	Piccolo 1, Flute 3, Flute 2, Flute 1, Flute 4, Violin 1	Piccolo 1, Flute 3, Flute 2, Flute 1, Flute 4, Violin 1
2. (green)	Horn 4, Horn 3	Horn 4, Horn 3
3. (purple)	Horn 2, Clarinet 1, Trumpet 2, Horn 1, Oboe 1, Clarinet 2, Bassoon 2, Trumpet 1, Oboe 3, Oboe 2	Horn 2, Clarinet 1, Trumpet 2, Horn 1, Oboe 1, Clarinet 2, Bassoon 2, Trumpet 1, Oboe 3, Oboe 2

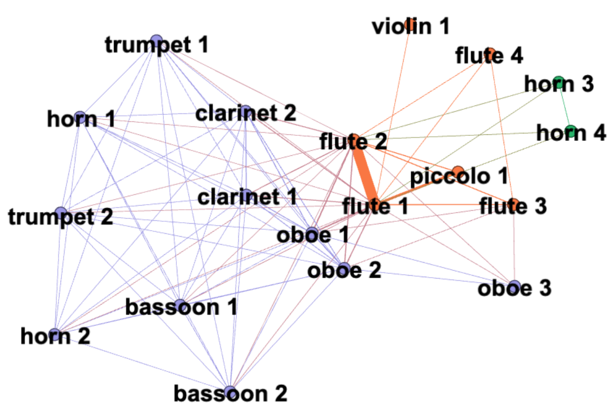

Figure 4.14 Flute Dominating Instruments Network with Communities Formed by the CNM Algorithm and the Louvain Algorithm

Unlike the Violin Dominating Instruments Network, this case-study instrument's dominating instruments network has enough nodes to have a more interesting structure. The communities formed by both algorithms are identical, each resulting in a modularity of 0.471. Although the modularity scores still indicate suboptimal partitioning, each group has a stronger identity that connects all of the instrument parts. The first group consists of flutes, piccolo, and the first violin part. These are instruments which frequently pair with the flute when it is dominating, either because they are mostly dominating instruments, like violin 1, or because they are of the flute family. The second group consists of the third and fourth horn parts, which are not frequently used horn parts. Group 3 consists of other instruments which do not dominate with the flute as frequently.

The centrality results are of interest, as they deem oboes 1 and 2 as important nodes in the Flute Dominating Instruments Network. As discussed in previous chapters, the flute and oboe are a common pairing, but rarely dominate together. However, the degree centrality and PageRank centrality seem to be highlighting the importance of this combination.

Table 4.8 reports the centrality results and the partitions produced for the Flute Embellishing Instruments Network. Figures 4.15 and 4.16 show the network visualizations produced by *Gephi*.

Table 4.8: Centrality & Community Detection Results for the Flute Embellishing Instruments Network

Centrality Results		
Rank	Degree Centrality	PageRank Centrality
1.	Oboe 1 (0.951)	Clarinet 1 (0.085)
2.	Clarinet 1 (0.927)	Clarinet 2 (0.083)
3.	Clarinet 2 (0.927)	Oboe 1 (0.079)
4.	Bassoon 1 (0.854)	Oboe 2 (0.070)
5.	Oboe 2 (0.854)	Bassoon 1 (0.063)
Community Results		
Group Number	Louvain Communities	CNM Communities
1. (green)	Xylophone, Violin 2, Viola, Cello, Bassoon 3, Flute 2, Harp 2, Violin 1, Contrabassoon, Bass, Harp 1	Xylophone, Violin 2, Viola, Cello, Bassoon 3, Flute 2, Harp 2, Violin 1, Contrabassoon, Bass, Harp 1
2. (purple)	Cornet 2, Cornet 1, Trombone 1, Horn 2, Alto Trombone, Trombone 2, Piccolo 1, Trumpet 2, Horn 1, Timpani 1, Horn 4, Triangle, Horn 3, Bass Trombone, Trombone 3, Trumpet 1	Trombone 1, Horn 2, Alto Trombone, Trombone 2, Piccolo 1, Trumpet 2, Horn 1, Timpani 1, Horn 4, Triangle, Horn 3, Bass Trombone, Trombone 3, Trumpet 1
3. (orange)	English Horn, Clarinet 3, Clarinet 1, Piccolo Clarinet, Flute 3, Oboe 4, Glockenspiel, Oboe 1, Flute 1, Bass Horn 2, Clarinet 2, Bassoon 2, Bassoon 1, Oboe 3, Oboe 2	Cornet 2, Cornet 1, English Horn, Clarinet 3, Clarinet 1, Piccolo Clarinet, Flute 3, Oboe 4, Glockenspiel, Oboe 1, Flute 1, Bass Horn 2, Clarinet 2, Bassoon 2, Bassoon 1, Oboe 3, Oboe 2

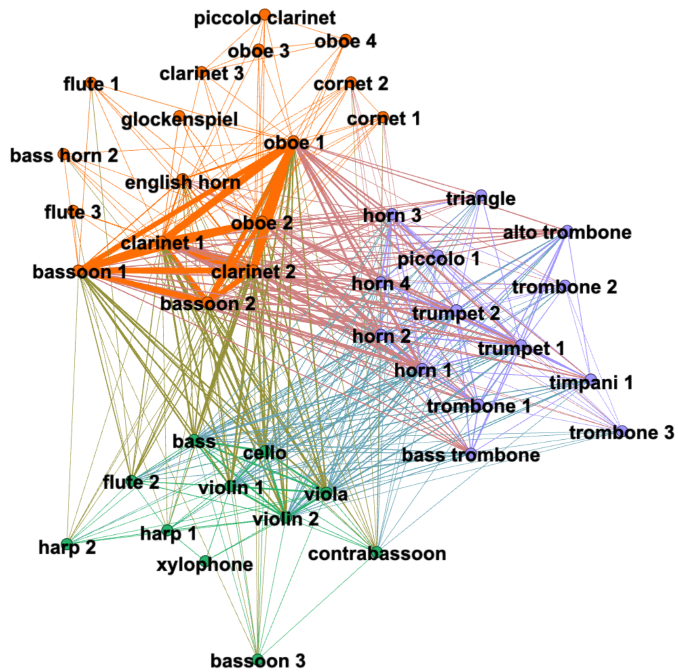

Figure 4.15 Flute Embellishing Instruments Network with Communities Formed by the CNM Algorithm

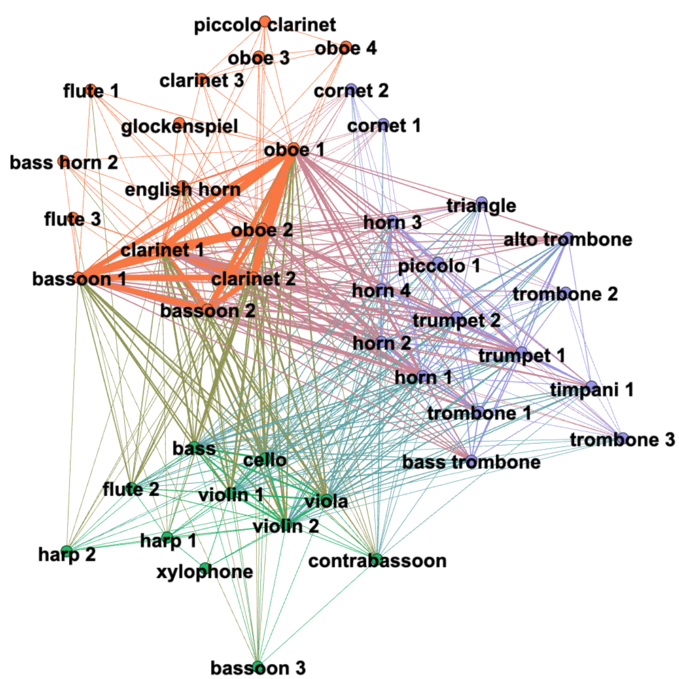

Figure 4.16 Flute Embellishing Instruments Network with Communities Formed by the Louvain Algorithm

Although the network visualizations for the flute's embellishing instruments network are similar to other embellishing instrument networks, its communities appear to be more separate from one another. There is no concentration of instrument nodes at the network's centre. This separation is reflected in the community detection's modularity, which sits at 0.303 for both sets of partitions. Although these values are still suboptimal, they indicate slightly stronger identities in the communities. The instruments in each group are similar, with both cornet parts being placed in group 2 for the Louvain communities and group 3 for the CNM communities. The first group consists primarily of string instruments, as well as some additional less frequently used instrument parts. The second group consists mainly of brass instruments. The third and final community consists of woodwind instruments. The family division is much more apparent in this network, yet it does not achieve the same level of optimal partitioning as the Dominating Instruments Network. This is due to the additional instruments added to the partitions, which impact the clarity of these communities' identities.

The centrality results indicate that the clarinets are the most important instrument nodes in this embellishing instruments network, in addition to having the highest degree and the heaviest edge. The oboe and the bassoon also received high rankings in both centrality algorithms. This is all in agreement with what was discovered in previous analyses of this data using relative frequencies and frequent pattern mining.

Oboe

Table 4.9 contains the most important nodes determined by the centrality algorithms, as well as the communities detected, in the Oboe Dominating Instruments Network. The visual representation of this network can be found in Figure 4.17.

Table 4.9: Centrality & Community Detection Results for the Oboe Dominating Instruments Network

Centrality Results		
Rank	Degree Centrality	PageRank Centrality
1.	Oboe 1 (1.000)	Oboe 1 (0.213)
2.	Oboe 2 (1.000)	Oboe 2 (0.195)
3.	Flute 1 (0.929)	Flute 1 (0.095)
4.	Flute 2 (0.929)	Flute 2 (0.094)
5.	Horn 1 (0.786)	Clarinet 1 (0.064)
Community Results		
Group Number	Louvain Communities	CNM Communities
1. (orange)	Clarinet 1, Flute 3, Flute 2, Flute 1, Clarinet 2	Clarinet 1, Flute 3, Flute 2, Flute 1, Clarinet 2
2. (purple)	Horn 2, Trumpet 2, Horn 1, Bassoon 2, Bassoon 1, Trumpet 1	Horn 2, Trumpet 2, Horn 1, Bassoon 2, Bassoon 1, Trumpet 1
3. (green)	English Horn, Oboe 1, Oboe 3, Oboe 2	English Horn, Oboe 1, Oboe 3, Oboe 2

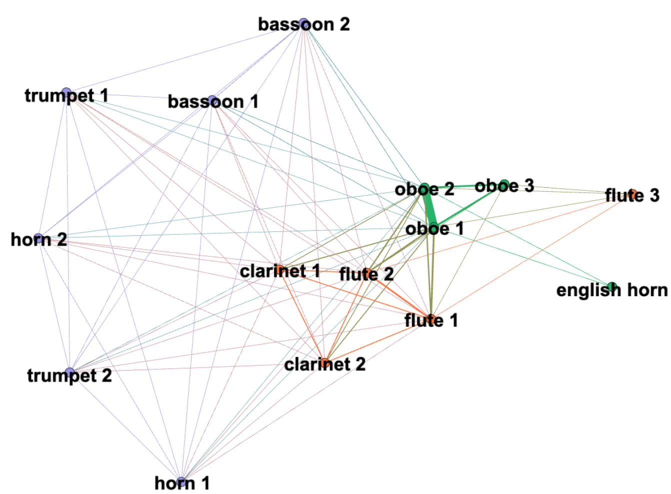

Figure 4.17 Oboe Dominating Instruments Network with Communities Formed by the CNM Algorithm and the Louvain Algorithm

The Oboe Dominating Instruments Network has a similar spatial structure to that of the flute. There is a clear separation between groups 2 and 3. However, there is no clear separation between groups 1 and 3. Both the Louvain and CNM algorithms produce the same partitions, achieving a modularity of 0.382. Although the partitions in this network are suboptimal, there are some common traits that can be observed among the different nodes in each group. Group 3 consists of the oboe parts and the English horn, which are all part of the oboe family. Group 1 consists of clarinet and flute parts, which are both woodwind instruments. As group 3 and group 1 both consist of woodwind instruments, this could explain why there is no explicit separation between the two groups in Figure 4.17. Group 2 contains more variety, as it consists of horn, trumpet, and bassoon parts. The previous two groups have timbral similarities or consist of one instrument family, but group 2 appears to be a catch-all for the other instrument parts that do not occur as frequently. This could be a contributing factor to its weak modularity score.

The most important nodes of this network are similar to the results from the Flute Dominating Instrument Network. Both algorithms determined that the oboes and flutes were some of the most important instrument parts in this case-study data. Once again, despite the flute not occurring the most frequently, these algorithms seem to have detected their strong partnership, as outlined in orchestration treatises.

Table 4.10 outlines the most important nodes and the partitions formed in the Oboe Embellishing Instruments Network. Figures 4.18 and 4.19 are visual representations of this graph in which the nodes are colour-coded based on their assigned communities.

Table 4.10: Centrality & Community Detection Results for the Oboe Embellishing Instruments Network

Centrality Results		
Rank	Degree Centrality	PageRank Centrality
1.	Clarinet 1 (0.921)	Clarinet 1 (0.092)
2.	Clarinet 2 (0.895)	Clarinet 2 (0.088)
3.	Flute 1 (0.789)	Flute 1 (0.070)
4.	Bassoon 2 (0.789)	Bassoon 1 (0.067)
5.	Bassoon 1 (0.789)	Bassoon 2 (0.064)
Community Results		
Group Number	Louvain Communities	CNM Communities
1. (orange)	Trombone 1, Horn 2, Trombone 2, Horn 1, Horn 4, Contrabassoon, Horn 3, Trombone 3, Oboe 2	Trombone 1, Horn 2, Trombone 2, Horn 1, Horn 4, Contrabassoon, Horn 3, Trombone 3, Oboe 2
2. (blue)	Trumpet 3, Trumpet 4, Trumpet 2, Timpani 1, Trumpet 1	Trumpet 3, Trumpet 4, Trumpet 2, Timpani 1, Trumpet 1
3. (green)	Xylophone, Piccolo Clarinet, Piccolo 1, Violin 2, Viola, Cello, Solo Violin, Harp 2, Violin 1, Bass, Harp 1	Xylophone, Piccolo Clarinet, Piccolo 1, Violin 2, Viola, Cello, Solo Violin, Harp 2, Violin 1, Bass, Harp 1, Clarinet 3, Flute 3
4. (pink)	Clarinet 3, Flute 3, English Horn, Clarinet 1, Alto Trombone, Bassoon 3, Bass Clarinet, Bass Horn 1, Flute 2, Flute 1, Bass Horn 2, Clarinet 2, Bassoon 2, Bassoon 1	English Horn, Clarinet 1, Alto Trombone, Bassoon 3, Bass Clarinet, Bass Horn 1, Flute 2, Flute 1, Bass Horn 2, Clarinet 2, Bassoon 2, Bassoon 1

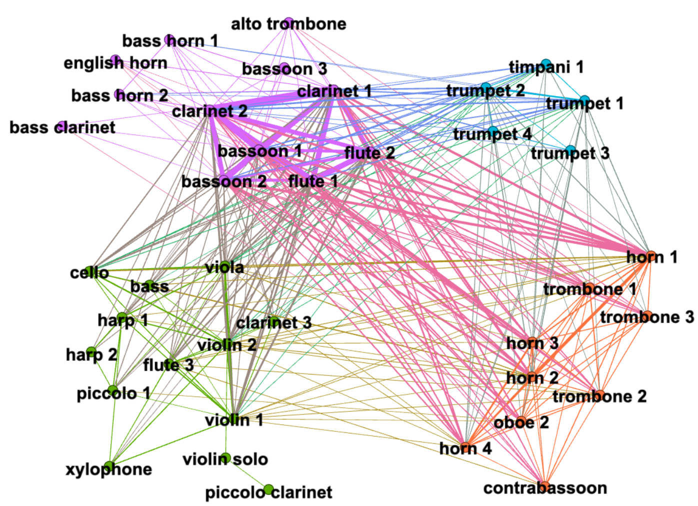

Figure 4.18 Oboe Embellishing Instruments Network with Communities Formed by the CNM Algorithm

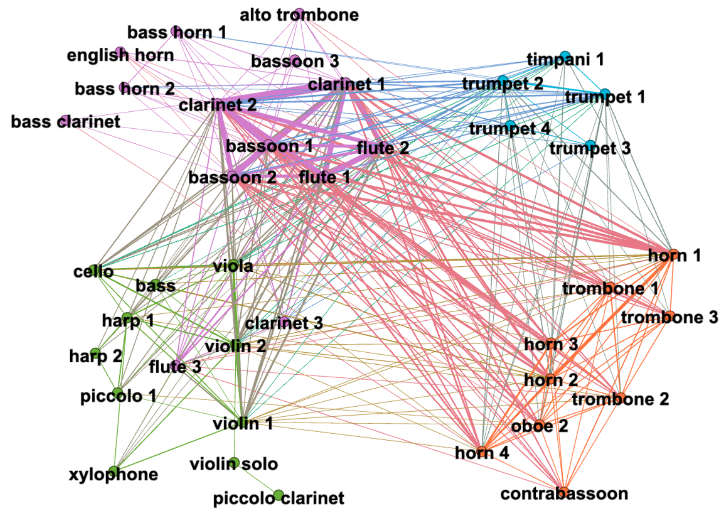

Figure 4.19 Oboe Embellishing Instruments Network with Communities Formed by the Louvain Algorithm

The figures show that *ForceAtlas* somewhat separated the communities. The CNM algorithm is more successful at this, as some members of group 4 are found near the group 3 instruments in Figure 4.17. It is also less compact than the flute's embellishing instruments network. This suggests that there is greater connection among the oboe's embellishing instruments than among the flute's embellishing instruments, as *ForceAtlas* repels less connected nodes.

Both centrality algorithms generated identical sets of important nodes, with slightly different rankings. Table 4.9 shows that the first oboe is an important node in the Flute Embellishing Instruments Network, but the reverse is not true for the Oboe Embellishing Instruments network (Table 4.10). The first two parts of the clarinet are much more important to the oboe network according to both measurement techniques. In this case, the order of the important instruments is identical to the order of the embellishing instruments in the distribution graph seen in Chapter 2's Figure 2.20.

The formed communities are, once again, suboptimal partitions, achieving a modularity of 0.317 and 0.306 for the CNM and Louvain algorithms, respectively. Some of the groups showed some cohesion among their members, consisting largely of one instrument family, such as the brass instruments in group 1 or the string instruments in group 3, or one type of instrument, like the trumpets in group 2. However, as is the case with many embellishing instruments, these groups have additional instruments that add noise to these partitions. Group 4, although it contains the heaviest edge and the highest degree node, has the greatest variety of instruments.

French horn

Table 4.11 shows the results for the centrality measurements and community detection conducted on the French Horn Dominating Instruments Network. The network's *Gephi* visualization is found in Figure 4.20.

Table 4.11: Centrality & Community Detection Results for the French Horn Dominating Instruments Network

Centrality Results		
Rank	Degree Centrality	PageRank Centrality
1.	Horn 1 (0.964)	Horn 1 (0.141)
2.	Horn 2 (0.964)	Horn 2 (0.140)
3.	Trumpet 2 (0.786)	Horn 3 (0.091)
4.	Trumpet 1 (0.786)	Horn 4 (0.087)
5.	Horn 4 (0.679)	Trumpet 1 (0.071)
Community Results		
Group Number	Louvain Communities	CNM Communities
1. (orange)	Horn 2, Horn 5, Horn 1, Horn 6, Horn 7, Horn 3	Horn 2, Horn 5, Horn 1, Horn 6, Horn 7, Horn 3
2. (blue)	Cornet 2, Cornet 1, Euphonium 2	Cornet 2, Cornet 1, Euphonium 2
3. (pink)	Clarinet 1, Cello, Oboe 1, Flute 2, Flute 1, Clarinet 2, Bass, Bassoon 2, Bassoon 1, Oboe 2	Clarinet 1, Cello, Oboe 1, Flute 2, Flute 1, Clarinet 2, Bass, Bassoon 2, Bassoon 1, Oboe 2
4. (green)	Trombone 1, Trumpet 3, Alto Trombone, Trombone 2, Trumpet 2, Tuba 1, Bass Trombone, Trombone 3, Trumpet 1	Trombone 1, Trumpet 3, Alto Trombone, Trombone 2, Trumpet 2, Tuba 1, Bass Trombone, Trombone 3, Trumpet 1

Of the four case-study instruments, the French horn has the greatest number of nodes in its dominating instruments network. Its communities also possess the lowest modularity scores, excluding the Violin Dominating Instruments Network, which had too few nodes to produce partitions. The CNM and Louvain algorithms both produced partitions with a modularity of 0.359. The nodes in each group have common traits when examining them. The first group contains all of the horn parts. The second group, the smallest, contains cornets and euphonium, which are rarely used. Group 3 consists of different woodwind instrument parts, as well as two low string instruments. The fourth group contains the other brass instruments. The fourth group contains the other brass instruments. Groups 1, 2, and 4 have very strong identities and all consist

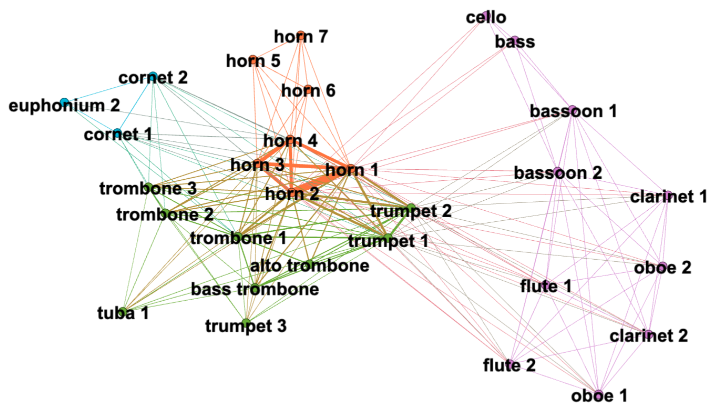

Figure 4.20 French Horn Dominating Instruments Network with Communities Formed by the CNM Algorithm and the Louvain Algorithm

of brass instruments. Group 3 is spatially separate from the other groups, which are much closer to one another. This could be due to group 3's inclusion of non-brass instruments. This varied group could also explain why the horn has the lowest modularity.

The centrality results are saturated with different horn parts, along with some trumpet parts. This contrasts with the other case-study instruments' most important nodes, which had more variety outside the case-study instrument. The results here reflect the horn's results in its relative frequency analysis, in which it was reported that it is paired with another horn part at a frequency exceeding 0.90. As such, it is no surprise that horn nodes are more important than any other instrument node in this network.

Table 4.12 details the French Horn Embellishing Instruments Network's centrality analysis and community detection results. The partitions, as well as the network itself, are visually represented in Figures 4.21 and 4.22.

Table 4.12: Centrality & Community Detection Results for the French Horn Embellishing Instruments Network

Centrality Results		
Rank	Degree Centrality	PageRank Centrality
1.	Bassoon 2 (0.978)	Bassoon 1 (0.073)
2.	Bassoon 1 (0.978)	Bassoon 2 (0.073)
3.	Clarinet 1 (0.935)	Clarinet 1 (0.070)
4.	Clarinet 2 (0.935)	Clarinet 2 (0.070)
5.	Oboe 1 (0.891)	Oboe 1 (0.061)
Community Results		
Group Number	Louvain Communities	CNM Communities
1. (orange)	Tubular Bells, Glockenspiel, Gong, Harp 2	Tubular Bells, Glockenspiel, Gong, Harp 1, Bass drum, Harp 2
2. (blue)	Trombone 1, Alto Trombone, Bass Trombone	Trombone 1, Alto Trombone, Bass Trombone, Violin 2, Viola, Cornet 2, Cello, Bass Clarinet, Timpani 1, Tuba 1, Violin 1, Contrabassoon, Cornet 1, Bass
3. (green)	English Horn, Clarinet 3, Clarinet 1, Piccolo Clarinet, Piccolo 1, Flute 3, Oboe 4, Trumpet 2, Oboe 1, Flute 2, Flute 1, Piccolo 2, Triangle, Clarinet 2, Cymbals, Trumpet 1, Oboe 3, Oboe 2	English Horn, Clarinet 3, Clarinet 1, Piccolo Clarinet, Piccolo 1, Flute 3, Oboe 4, Trumpet 2, Oboe 1, Flute 2, Flute 1, Piccolo 2, Triangle, Clarinet 2, Cymbals, Trumpet 1, Oboe 3, Oboe 2, Horn 2, Trombone 2, Bassoon 3, Horn 1, Horn 4, Bassoon 2, Horn 3, Trombone 3, Bassoon 1
4. (pink)	Horn 2, Trombone 2, Bassoon 3, Horn 1, Horn 4, Bassoon 2, Horn 3, Trombone 3, Bassoon 1, Violin 2, Viola, Cornet 2, Cello, Bass Clarinet, Timpani 1, Tuba 1, Violin 1, Contrabassoon, Cornet 1, Bass, Bass Drum, Harp 1	N/A

Although the French Horn Embellishing Instruments Network does not have the highest number of edges of the case-study embellishing instrument graphs, its edges are very thick, especially between groups 2 and 3 in Figure 4.21. *ForceAtlas* also formed distinct separations between the CNM groups, as seen by this figure. In Figure 4.22, on the other hand, the groups appear to be mixed together. It should be noted that the Louvain algorithm produced an additional community, whereas the CNM algorithm created only three partitions.

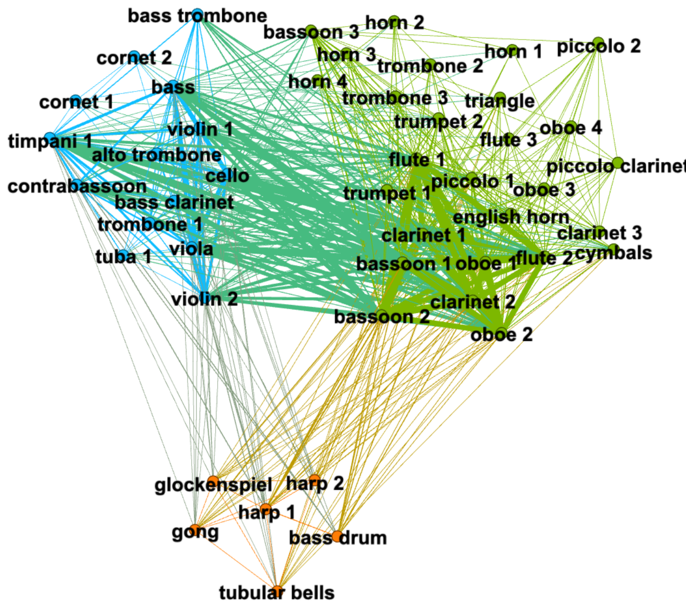

Figure 4.21 French Horn Embellishing Instruments Network with Communities Formed by the CNM Algorithm

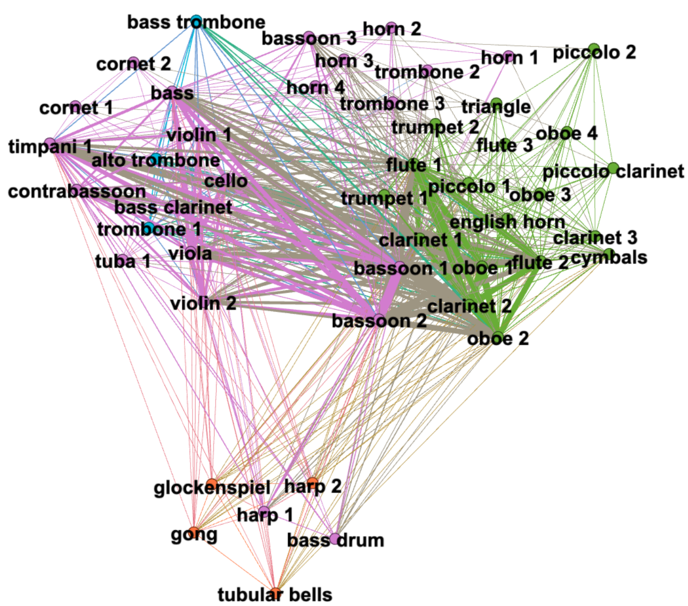

Figure 4.22 French Horn Embellishing Instruments Network with Communities Formed by the Louvain Algorithm

The modularity of the Louvain algorithm and the CNM algorithm's groupings are 0.297 and 0.328, respectively. Once again, the modularity algorithm used favours larger groups, which is why the Louvain communities scored a lower modularity than the CNM partitions. The Louvain algorithm extracted an additional group of trombones, and formed a fourth group using nodes found in groups 1, 2, and 3 formed by the CNM algorithm. Group 3 in both algorithms grouped together some of the woodwind and brass instruments with a higher pitch range, as well as the cymbals. The fourth group in the Louvain algorithm and the second group in the CNM algorithm consist of strings, brass, and woodwinds with lower pitch-ranges and similar timbres. The French Horn Embellishing Instruments Network falls into the common trend of these embellishing networks. The slight structure to these communities is muddled by additional less frequently occurring instrument nodes, which in turn impacts the modularity results.

The most important nodes, according to the two algorithms, are the bassoon, the clarinet, and the oboe. Once again, this mirrors the results in the relative frequency distribution. The flute, which has appeared in many of these tables, is notably absent in Table 4.12. This could be due to the stark difference between the French horn's and flute's timbres, which makes them less likely to be paired in a dominating-embellishing relationship.

4.4 Discussion

The network analysis conducted in this chapter explored the instrumentation and the potential subgroupings in both timbral augmentation blend roles. Both community detection algorithms produced the same sets of partitions across each of the dominating instrument networks. The communities detected in these networks were also stronger, especially the Dominating Instruments Network, which achieved the highest modularity of any network present in this analysis. Chapters 2 and 3 discussed the fact that dominating instrument groups often consist of only one instrument family and rarely use multiple families at a time. As such, the networks produced smaller and more compact partitions with strong links. The embellishing instruments networks, on the other hand, have a more web-like structure and their communities' common traits are more difficult to decipher. This can be traced back to the fact that embellishing instruments blend with each other and have much larger groups. This promotes a greater variety of instrument parts in an excerpt's embellishing instrument groups. Rather than having a restriction to one family of nearly identical timbres, the restrictions are much looser. This also makes it much more difficult to analyze the groups and find common factors among the different instrument nodes. Thus, from these results, there are not any detectable subgroups within the embellishing instrument groups in the same way that they exist in the dominating instruments.

The expectation for the community detection algorithm was to detect functional subcategories within the blend roles, particularly the embellishing instruments group. Some of the underlying community structures were interesting, despite their suboptimal modularity. The embellishing instrument networks, in particular, explored more interesting common factors beyond the instrument's family. Some groups were formed based on the frequency at which its instrument nodes occurred in the data. For instance, the Embellishing Instruments Network had some communities that consisted of less frequently occurring instrument nodes, like group 1's wind instruments or group 2's various percussion instruments. The embellishing instrument networks in the case-study instruments found different ways to partition nodes, and created groups of instruments with similar timbral qualities or similar pitch ranges. Although timbre and pitch range were not included in any of OrchARD's annotations, the community detection algorithms still grouped these instruments together based on these similarities. From these results alone, it remains difficult to determine if this was coincidental or if the CNM and Louvain algorithms were able to detect these specific parameters.

The Timbral Augmentation Network was not as insightful as the other ten networks. The centrality results were interesting, as each algorithm focused on a different blend role. However, the communities produced by one algorithm were vastly different from those produced by the other detection algorithm. In addition, it was more challenging to determine how the communities were formed. What did the instrument parts in a given community have in common with one another? Their shared traits were more unclear than was the case with the embellishing instruments network. The expectation when creating this network was that there were some instruments that were more strictly dominating instruments and others that were strictly embellishing instruments. However, in reality, few instruments adhere to that absolute categorization. The context of how many instruments are involved, as well as which instruments are involved, can influence which blend role that an instrument falls into. This can even be seen when comparing the different results across tables. Table 4.6 ranks oboe 1 as the most important embellishing instrument for the violin case study, but the same instrument part is considered to be an important dominating instrument for flutes in Table 4.7. As many instruments can be either dominating or embellishing, depending on which instruments are involved in the blend, it adds a level of complexity which cannot be accurately portrayed by this network and the community detection algorithms applied to it. The Timbral Augmentation Network, although an interesting concept in theory, removes a lot of the context that is imperative when investigating a timbral augmentation blend.

Chapters 2 and 3 discussed how certain instruments or instrument families fall into specific blend roles. String instruments dominate more frequently. Woodwind instruments are more likely to be used as embellishing instruments. Brass instruments have more versatility and can be used in different blend roles at a nearly equal rate. Percussion are rarely used, but are more prominently used in embellishing instrument groupings. These findings are reflected in some of the results of this network analysis. The Dominating Instruments Network ranked string instruments, specifically violin, as the heaviest edges and the nodes with the highest degree. This supports the Read (1979) claim that the strings are the most frequently used family in Classical compositions. In contrast, brass instruments were ranked as the most important instrument nodes, which confirms their versatility as dominating instruments. That is, brass instrument nodes share edges with a greater variety of instrument nodes as opposed to string instruments. This mirrors the claim made in the Adler (1982) orchestration treatise regarding the versatility of certain brass instruments due to their tone quality, particularly the French Horn. In other words, brass instruments, like the French horn, can be paired with a large variety of instruments and sound pleasant in the standards of Western music. Woodwinds appear as the most important nodes in every embellishing instrument network. Percussion instruments are frequently found on the outer edge of embellishing instrument networks, indicating that they are rarely used. These observations from the networks confirm what has been concluded by the relative frequency analysis and the frequent pattern mining conducted in the previous two chapters.

The flute and oboe pairing returns once again in this analysis (see sections 2.2.2.3 and 3.3.2.3). The relative frequency analysis and the frequent pattern mining showed that this pairing occurs more often as a dominating-embellishing relationship. Despite this, the oboe was deemed an important node by the centrality algorithms in the Flute Dominating Instruments Network and vice versa. This is another instance of the network analysis detecting relationships beyond their relative frequency of occurrence. Even if their relative frequency of occurrence as two dominating instruments or as a dominating-embellishing instrument relationship does not highlight the importance of their relationship, the centrality algorithms were able to successfully detect and highlight this relationship.

Networks have proven to be an insightful tool for analyzing blends. The different algorithms provided greater insight to OrchARD's data. The centrality algorithms were able to determine which instruments were the most versatile. The results of these algorithms created a point of comparison with the heaviest edge or the highest degree, both of which relied more on relative frequency or blend strength. The community detection algorithms, even when producing weak partitions, were able to group instruments with timbral or pitch similarities. Unfortunately, the results from these algorithms did not detect strong enough subgroupings for different instruments in each blend role.

The visual representations added a dimension to the analysis of timbral augmentations. For instance, they were able to give insight to the modularity scores, or even show the difference between how the two community detection algorithms partitioned instruments with respect to the groups' internal connectedness. They also managed to highlight patterns beyond the limited data found in each annotation. The frequent pattern mining in Chapter 3 introduced larger groupings beyond the groups of two provided by the relative frequencies of occurrence analysis in Chapter 2. These larger common groupings are highlighted by the thick edges in certain networks. For instance, the Flute Embellishing Instruments Network (Figures 4.15 and 4.16) has a triangle of heavy edges between oboe, bassoon, and clarinet parts. The blend strength of these instruments and the frequency at which they occur together is sharply contrasted by the weaker links between other instruments. For computational purposes, Chapter 3's frequent pattern mining and this chapter's centrality metrics were limited to a certain number of results. The network visualizations not only confirm previous results, but they also give a greater frame of reference. Additionally, the edges between different communities, such as the ones in Figures 4.21 and 4.22, display the numerous embellishing instrument combinations that exist, whether they occur frequently or not. Ultimately, the networks were able to successfully create a more detailed image of how dominating and embellishing instruments interact with each other, especially within their respective blend roles.

Chapter 5

Conclusion

5.1 General discussion

The current study sought to analyze orchestration patterns found in OrchARD's timbral augmentation blend data to explore how instruments are placed into dominating or embellishing blend roles. By using three different analysis techniques, the data were able to be examined with different lenses and patterns in the assignment of instruments to certain blend roles were able to be identified.

The initial analysis in Chapter 2 calculated the relative frequencies of occurrence for each instrument and instrument part to determine early patterns in each blend role. This chapter also incorporated the timbral emergence data as a way to compare the embellishing and emerging instruments. The relative frequency analysis formed simple associations between instrument families and certain blend roles, and even highlighted frequently used instrument families for each blend role. Strings, particularly the violin, were associated with dominating instruments. Woodwinds, like the clarinet and the bassoon, were more likely to be used as embellishing or emerging instruments. These two assignments confirmed the initial hypothesis (see section 1.3.1), as strings were often given the melodic line and the woodwinds enhanced them. Brass instruments, especially the French horn, had greater versatility, meaning that they were frequently used in all blend roles. The case-study instruments gave further insight into some of the patterns in the orchestration of these blend roles. In particular, it was found that instruments from the same family are more likely to dominate at the same time (see section 2.2.2). For instrument families that included a greater number of instruments, such as the woodwinds, dominating instrument groups comprised woodwind instruments with similar timbres, such as the flute and the piccolo. The relative frequency analysis also indicated some key differences between embellishing and emerging instruments, namely the impact that the different group sizes had on the instruments selected for the respective blend role (see section 2.3). These findings provided basic insight into OrchARD's blend contents and built a foundation for the subsequent complex analysis techniques.

Chapter 3 further explored the patterns of instruments used alone or conjointly in different blend roles seen in Chapter 2 using a frequent pattern mining algorithm. This analysis technique allowed for the group sizes of instrument patterns to be expanded beyond the pairings interpreted from the relative frequency results. In doing so, some of the smaller patterns identified in Chapter 2, like the flute and oboe pairing in a dominating-embellishing relationship or as emerging instruments, were expanded and included other instruments, like the clarinet. Other instruments' proclivity for certain blend roles or pairings with certain instruments were further confirmed. This included the string family's tendency to be dominating instruments and the bassoon's tendency to embellish the French horn. Findings such as these gave further depth to the patterns discovered solely using the relative frequencies of occurrence. In addition, the association rule metrics of confidence, lift, leverage, and conviction provided a way to examine the strength of an association rule and the statistical dependence of its items. Ultimately, it was noticed that cases of instrument part doubling resulted in greater statistical dependence than cases of two different instruments playing together (see section 3.4).

Chapter 3 also explored the impact of blend strength on the generated frequent patterns. However, upon adjusting the blend strength threshold in the current study's variation of the FP-Growth algorithm, there was no change to the generated itemsets beyond minor adjustments to the association rule metrics. Of course, if the blend strength was set to a high value, no frequent itemsets were detected. This is likely due to the normal distribution of the blend strength distribution (see section 2.1.1). If an itemset occurred frequently enough to be detected by the FP-Growth algorithm, then its instances of weaker perceived blend strength occurred as often as its instances of very high perceived blend strength.

The network analysis conducted in Chapter 4 sought to focus on the relationships between the instrument parts in different timbral augmentation blend roles. The network visualizations gave further insight to how the different instrument parts interact with one another beyond the most frequently occurring instruments or instrument groupings. For instance, the Dominating Instruments Network had a separate component consisting of both harp parts. Although harps are rarely used in timbral augmentation blends, the visual representations and the network statistics were able to identify this unique phenomenon, which was overlooked in the previous statistical analyses. The different shapes produced by Gephi also revealed an interesting point of comparison between the dominating and the embellishing instrument networks. The dominating instrument networks had a clearer separation between the different instrument groups, whereas the embellishing instrument networks adopted a web-like structure. This difference provided a visual confirmation to the previously explored differences between the orchestrations of the two blend roles.

Chapter 4 also explored potential subgroupings within the blend roles and identified their important instruments using community detection and centrality, respectively. The community detection algorithms identified stronger communities within the dominating instrument networks, but were less successful in doing so for the embellishing instrument networks. This may be the result of multiple instruments across instrument families blending, meaning that any distinction or function beyond the blended sound are not considered. However, the embellishing instrument networks' communities appear to have been formed by subtler similarities between the instruments. For instance, some communities consisted of instruments with similar pitch ranges or similar relative frequencies of occurrence (see sections 4.3.1.3 and 4.3.2.4). The centrality algorithms also highlighted the difference between high centrality scores and high relative frequency of occurrence: versatility. An instrument part that occurs frequently in the data does not necessarily make it an important instrument according to the centrality algorithms. For example, in the Dominating Instruments Network, although the violin parts occurred the most in the data, brass instrument parts were deemed to be more important by both degree centrality and PageRank centrality. Overall, the combination of visual elements and statistical measurements in network analysis gave further insight to the previous results and explored new angles to the data that were not computationally possible with the previous analysis techniques.

The results across all three chapters share similarities with what had been previously written in 19th and 20th century orchestration treatises. The hypothesis stated in Chapter 1 (see section 1.3.1) theorized that instruments that took on melodic lines more frequently would also dominate in timbral augmentations more frequently. Read (1979) noted that strings, especially in the Classical repertoire, were given the melodic line more frequently. Piston (1955) also noted that woodwinds had a greater number of weak points in their range compared to other families, and were often used to strengthen the melodic line. This is reflected perfectly in the strings' tendency to dominate and the woodwinds' tendency to embellish in timbral augmentation blends. The hypothesis is therefore supported by these analyses.

The preference to use certain instruments in one blend role over another was addressed in these writings, albeit using different terminology. For instance, the oboe's greater use as an emerging instrument was mentioned by Berlioz (1948), who wrote about its preferred placement in a large ensemble section so as to make its sound less distinct. In addition, these treatises indirectly addressed the versatility of the brass family, as it was seen to be used frequently across all blend roles examined. Adler (1982) explored this phenomenon in the French horn specifically, crediting the instrument's tone for its ability to complement many different instrument timbres. Although the terminology and methodology used to create these annotations differs from the terms used in orchestration treatises, this corpus analysis managed to arrive at similar conclusions for blend roles.

The flute-oboe relationship is mentioned heavily in orchestration treatises, and is seen in the results across the different analysis techniques. The current study, however, provided clarifications of this relationship. Adler (1982) cited the interesting timbre produced by the flute and oboe combination and tracked its use across multiple works, including Schubert's *Symphony 8* and Debussy's *La Mer*. In the current study, this pairing seemed to be most prominent in timbral emergence blends. Although this pairing exists in timbral augmentation blends, the flute and the oboe rarely dominate at the same time (see sections 2.2.2.2 and 2.2.2.3). Rather, when one of the two instruments dominates, the other is relegated to an embellishing role. It is in the timbral emergence excerpts that their blended timbre is most often seen (see section 2.3). Adler uses the terms *neutralization* and *balance* when describing the resulting tone of the flute and the oboe playing together. Through this analysis, this relationship was clarified to be preferable in blend roles that require all the instruments involved to blend together. It is not entirely surprising, however, considering the language in orchestration treatises used to describe their combination.

The separation by instrument family within the dominating instrument groupings was seen in each chapter. These results are reminiscent of the reported results from Le et al. (2022), who investigated patterns in textural layers of orchestral pieces. Their concept of textural layers is similar to the McAdams et al. (2022) definition of sequential grouping. Le et al. found that a given layer, regardless of its functional role in the texture, rarely featured multiple instrument families at once. For example, if the strings were found in the melodic layer, it was rare for brass instruments to be in that same layer. They also examined how the instruments across different layers interact with and blend with one another. For instance, the brass section is perceived to blend more effectively with the woodwind section than with the string section.

Each analysis technique used in this corpus analysis confirmed a phenomenon similar to that in Le et al. within the dominating instruments for a timbral augmentation blend. The instrument family division by layer could be a result of how listeners group instruments together perceptually. The Gestalt principle of similarity states that similar sounds are grouped together and that these groupings are delimited by acoustic discontinuities (McAdams, 2019). This concept extends to include timbral similarities and differences. Upon inspecting timbral augmentation annotations found in OrchARD, it was noticed that those having multiple instrument timbres dominating at once shared some similarities with timbral heterogeneity. As explained in their notes left by annotators, this resulted in a weaker blend. From these annotations, it can be hypothesized that, on the concurrent grouping level, if multiple instrument timbres were to dominate at once in a timbral augmentation blend while remaining distinct from one another, the blend risks becoming a weak hybrid of a timbral augmentation and timbral heterogeneity. Therefore, in order to produce a stronger blend, it is preferable to have instruments of one or more similar timbres dominate the resulting blend. These similarities between the current study and Le et al.'s study demonstrate cohesion among the different auditory grouping processes.

5.2 Limitations

The current study contains a set of limitations, largely as a result of the OrchARD data. The OrchARD data were created using excerpts from Classical and Romantic orchestral pieces (see section 1.3.2). The lone exception, Ticheli's *Symphony 2*, was omitted from the corpus analysis as it was considered to be an outlier piece, having been composed for wind orchestra rather than symphonic orchestra. So, OrchARD's representation of different musical eras, genres, and instruments is greatly limited. For instance, most of the annotations which featured saxophones were from the Ticheli piece, which is understandable as the saxophone was invented much later than the other instruments found in OrchARD and are prevalent in concert band pieces. In removing that symphony, only a few instances of the first alto saxophone part remained. In addition, the number of composers is limited to 17 popular Western composers from the Classical and Romantic eras, and the number of annotations associated with each composer varies. For example, Borodin has only 17 annotations, whereas Debussy has 130. This phenomenon can also be a result of the different lengths of their two pieces. Borodin's *In the Steppes of Central Asia* is a symphonic poem, and Debussy's *La Mer* is a much larger symphonic work with three movements. Therefore, the difference in the number of annotations is partially a result of the musical reality. From a statistical standpoint, however, it becomes a limitation.

The different blend types are also not represented equally in the data. As mentioned in Chapter 2, there are 1,074 timbral augmentation annotations and 313 timbral emergence annotations. These counts exclude the annotations contributed by Ticheli's *Symphony 2*. Nevertheless, the different number of annotations poses an issue when accurately comparing the difference between blend roles, like the embellishing and emerging instruments. The greater number of timbral augmentation annotations allows for some outliers to appear, such as the annotations in which the harps dominate the blend. The timbral emergence annotations appear to mainly feature instances where the orchestra plays *in tutti*. There are some instances of timbral emergence blends with smaller groups of instruments; however, these cases are offset by the numerous instances with larger ensembles. If there were also 1,074 timbral emergence annotations, the results could be vastly different. The smaller number of timbral emergence blends may be reflective of the musical reality for the epochs analyzed, but it once again poses an issue for the corpus analysis. One possible solution to this issue is to seek out works that feature more timbral emergence blends than timbral augmentation blends, which poses its own set of challenges. The more realistic approach may be to select statistical tools that address the different sample sizes and to remain conscious of these differences when interpreting the results.

The annotations themselves are also limited in the amount of information provided about the recording and the score. Chapter 1 mentions how certain recording years are missing from the OrchARD documentation (see section 1.3.2). The annotations were done based on what the experts heard, and thus are highly dependent on the recording used. This missing information becomes an issue when outlier annotations or weak blends are considered. For instance, if a certain annotation appeared to be out of the ordinary, it may have been the result of a typo or an error in the annotation's entry in the database. Without the recording year, there is no way to verify that information. In addition, the annotations do not mention details regarding the performance techniques or the dynamics employed in the excerpt. Some of these details are mentioned in the notes section, but this is not the case for every annotation. Certain performance techniques, such as playing pizzicato on string instruments, can have great influence on the blend, particularly when juxtaposed with sustaining instrument sounds (see section 2.3). Different instrument dynamics can also have an impact. For instance, the woodwinds were able to dominate over the strings in Mahler's *Symphony 1* due to the woodwinds playing piano as the strings played triple piano. Although this information can be found in the score, it was not feasible to search for these details for every annotation used in this corpus analysis as they were not encoded into the annotation data.

5.3 Future directions

This corpus analysis was a first exploration of OrchARD contents at length. As such, certain details relating to timbral augmentation blends were overlooked and the analysis techniques were simplified to give an overview. The results from the current study are promising and have provided potential directions for future research. For instance, in the interest of continuing the investigation of patterns in timbral augmentation and timbral emergence blends, the next step would be to consider the differences between punctuated, stable, and transforming blends (see section 1.2.2). Are there more specific patterns in their orchestration than can be found in timbral augmentations and timbral emergences at large? This information is readily available in the OrchARD annotations, but it was not explored in the current study. Chapter 2 (see section 2.3) briefly touched on the difference between the instruments used in punctuated and sustained timbral emergence blends. The differences in the average number of instrument parts used, as well as the increased use of percussion instruments in punctuated timbral emergence blends, give strong reasons to explore punctuated, stable, and transforming blends further. If this information was uncovered in a brief glance, what other patterns are found in these blend subtypes?

For a potential follow-up study on transforming blends, it is imperative to use a technique that allows for temporal changes. Network analysis does allow for time-variation in the form of *temporal networks*. Temporal networks introduce time as a new variable for links. For instance, a link between nodes A and B may be active at one moment in time, but inactive the next. This makes temporal networks popular for tasks such as tracking epidemics in the medical field, and could make them useful for exploring transforming blends (Barabási, 2016). Two temporal networks can be constructed to show the evolution of certain instruments taking turns dominating the timbre in each transforming blend excerpt, and then serve as a basis for comparison and tracking patterns.

Temporal networks have an additional use in future analyses of OrchARD data. The current study considered each annotation independently of its movement, piece, composer, or year of composition. Temporal networks may be used to track the evolution of orchestration within any of these aforementioned categories. The results, of course, would be restricted to the information available in OrchARD at the time of the experiment. However, the findings could be used to answer specific questions about blend in a certain piece or era.

As previously explained, the limitations of the current study can be primarily traced back to the limited data in OrchARD. The database is still growing, as platforms like OrchView[1] allow for auditory grouping annotations to be exported to OrchARD. Indeed, these future annotations have the potential to either cement the findings of the current study or to refute them entirely. They could also clarify certain questions that still remain at the end of the current analysis, like the validity of the timbral emergence patterns. Many of these possibilities, however, are dependent on the method in which these new annotations are created and accepted. If the newer annotations are from pieces written by the same composers already featured in OrchARD, many of the current issues could persist. The scope of the data will remain limited. Ideally, annotations will be collected such that the variety in the composers and the number of annotations associated with each orchestral effect will be balanced. An increased number of annotations also opens the possibility of conducting machine learning experiments on the OrchARD data. That prospect, however, remains contingent on abiding by the aforementioned recommendations for future data collection.

[1]https://www.actorproject.org/workgroups/orchview

References

Adler, S. (1982). *The study of orchestration*. W. W. Norton & Company, Inc.

Aggarwal, C. C. (2015). Association pattern mining. In *Data mining: The textbook* (pp. 93–133). Springer International Publishing. https://doi.org/10.1007/978-3-319-14142-8_4

Agrawal, R., & Srikant, R. (1994). Fast algorithms for mining association rules. In *Proceedings of the 20th International Conference on Very Large Data Bases, (VLDB'94)* (pp. 487–499). h3ps://www.vldb.org/conf/1994/P487.PDF

Aguilar-Ruiz, J., Rodríguez -Baena, D., & Alves, R. (2013). Frequent pattern mining. In W. Dubitzky, O. Wolkenhauer, K.-H. Cho, & H. Yokota (Eds.), *Encyclopedia of systems biology* (pp. 760–761). Springer New York. https://doi.org/10.1007/978-1-4419-9863-7_1114

Antoine, A., Depalle, P., Macnab-Séguin, P., & McAdams, S. (2021). Modeling human experts' identification of orchestral blends using symbolic information. In R. Kronland-Martinet, S. Ystad, & M. Aramaki (Eds.), *Perception, representations, image, sound, music* (pp. 363–378). Springer International Publishing. https://doi.org/10.1007/978-3-030-70210-6_24

Barabási, A.-L. (2016). *Network science* (1st ed.). Cambridge University Press.

Berlioz, H. (1948). *Treatise on instrumentation* (R. Strauss, Ed.; T. Front, Trans.). Edwin F. Kalmus.

Bregman, A. S. (1990). *Auditory scene analysis: The perceptual organization of sound*. The MIT Press. https://doi.org/10.7551/mitpress/1486.001.0001

Cano, P., Celma, O., Koppenberger, M., & Buldú, J. M. (2006). Topology of music recommendation networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, *16*(1), 013107. https://doi.org/10.1063/1.2137622

Chon, S. H. (2022). Corpus approaches to orchestration. In D. Shanahan, J. A. Burgoyne, & I. Quinn (Eds.), *The Oxford handbook of music and corpus studies* (1st ed.). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780190945442.013.14

Chon, S. H., Huron, D., & DeVlieger, D. (2018). An exploratory study of Western orchestration: Patterns through history. *Empirical Musicology Review*, *12*(3), 116. https://doi.org/10.18061/emr.v12i3-4.5773

Clauset, A., Newman, M. E. J., & Moore, C. (2004). Finding community structure in very large networks. *Physical Review E*, *70*(6), 066111. https://doi.org/10.1103/PhysRevE.70.066111

Cormen, T. H., Leiserson, C. E., Rivest, R. L., & Stein, C. (2022). *Introduction to algorithms* (4th ed.). The MIT Press.

Gao, S., Ma, J., Chen, Z., Wang, G., & Xing, C. (2014). Ranking the spreading ability of nodes in complex networks based on local structure. *Physica A: Statistical Mechanics and its Applications*, *403*, 130–147. https://doi.org/10.1016/j.physa.2014.02.032

Geng, L., & Hamilton, H. J. (2006). Interestingness measures for data mining: A survey. *ACM Computing Surveys*, *38*(3), 9. https://doi.org/10.1145/1132960.1132963

Gleiser, P., & Danon, L. (2003). Community structure in jazz. *Advances in Complex Systems*, *06*(4), 565–573. https://doi.org/10.1142/S0219525903001067

Goodchild, M., & McAdams, S. (2018). Perceptual processes in orchestration. In E. I. Dolan & A. Rehding (Eds.), *The Oxford handbook of timbre* (pp. 496–524). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780190637224.013.10

Grahne, G., & Zhu, J. (2003). Efficiently using prefix-trees in mining frequent itemsets. *Proceedings of the International Workshop on Frequent Itemset Mining Implementations*, 123–132.

Grey, J. M. (1977). Multidimensional perceptual scaling of musical timbres. *The Journal of the Acoustical Society of America*, *61*(5), 1270–1277. https://doi.org/10.1121/1.381428

Han, J., Cheng, H., Xin, D., & Yan, X. (2007). Frequent pattern mining: Current status and future directions. *Data Mining and Knowledge Discovery*, *15*(1), 55–86. https://doi.org/10.1007/s10618-006-0059-1

Han, J., Pei, J., & Yin, Y. (2000). Mining frequent patterns without candidate generation. *ACM SIGMOD Record*, *29*(2), 1–12. https://doi.org/10.1145/335191.335372

Heaton, J. (2016). Comparing dataset characteristics that favor the Apriori, Eclat or FP-Growth frequent itemset mining algorithms. *SoutheastCon 2016*, 1–7. https://doi.org/10.1109/SECON.2016.7506659

Hegland, M. (2007). The Apriori algorithm: A tutorial. In S. S. Goh, A. Ron, & Z. Shen (Eds.), *Mathematics and computation in imaging science and information processing* (pp. 209–262). World Scientific. https://doi.org/10.1142/9789812709066_0006

Hevey, D. (2018). Network analysis: A brief overview and tutorial. *Health Psychology and Behavioral Medicine*, *6*(1), 301–328. https://doi.org/10.1080/21642850.2018.1521283

Johnson, R. (2011). The standard, power, and color model of instrument combination in Romantic-era symphonic works. *Empirical Musicology Review*, *6*(1), 2–19. https://doi.org/10.18061/1811/49759

Kendall, R. A., & Carterette, E. C. (1991). Perceptual scaling of simultaneous wind instrument timbres. *Music Perception*, *8*(4), 369–404. https://doi.org/10.2307/40285519

Kendall, R. A., & Carterette, E. C. (1993). Identification and blend of timbres as a basis for orchestration. *Contemporary Music Review*, *9*(1), 51–67. https://doi.org/10.1080/07494469300640341

Le, D.-V.-T., Giraud, M., Levé, F., & Maccarini, F. (2022, July 28). A corpus describing orchestral texture in first movements of Classical and early-Romantic symphonies. In *Proceedings of the 9th International Conference on Digital Libraries for Musicology* (pp. 27–35). ACM. https://doi.org/10.1145/3543882.3543884

McAdams, S. (2019). Timbre as a structuring force in music. In K. Siedenburg, C. Saitis, S. McAdams, A. N. Popper, & R. R. Fay (Eds.), *Timbre: Acoustics, perception, and cognition* (pp. 211–243). Springer International Publishing. https://doi.org/10.1007/978-3-030-14832-4_8

McAdams, S., & Bregman, A. (1979). Hearing musical streams. *Computer Music Journal*, *3*(4), 26–43+60. http://www.jstor.org/stable/4617866

McAdams, S., Goodchild, M., & Soden, K. (2022). A taxonomy of orchestral grouping effects derived from principles of auditory perception. *Music Theory Online*, *28*(3). https://doi.org/10.30535/mto.28.3.6

Nalousi, S., Farhang, Y., & Babazadeh Sangar, A. (2021). Weighted frequent itemset mining using weighted subtrees: WST-WFIM. *IEEE Canadian Journal of Electrical and Computer Engineering*, *44*(2), 206–215. https://doi.org/10.1109/ICJECE.2020.3035472

Needham, M., & Hodler, A. E. (2019). *Graph algorithms : Practical examples in Apache Spark and Neo4j*. O'Reilly.

Newman, M. E. J. (2018). *Networks* (2nd ed.). Oxford University Press.

Page, L., Brin, S., Motwani, R., & Winograd, T. (1999). *The PageRank citation ranking: Bringing order to the web.* (Technical Report). Stanford InfoLab. http://ilpubs.stanford.edu:8090/422/

Park, D., Bae, A., & Park, J. (2014). The network of Western Classical music composers. In P. Contucci, R. Menezes, A. Omicini, & J. Poncela-Casasnovas (Eds.), *Complex networks V* (pp. 1–12). Springer International Publishing. https://doi.org/10.1007/978-3-319-05401-8_1

Park, D., Bae, A., Schich, M., & Park, J. (2015). Topology and evolution of the network of Western Classical music composers. *EPJ Data Science*, *4*(1), 2. https://doi.org/10.1140/epjds/s13688-015-0039-z

Piston, W. (1955). *Orchestration* (1st). W. W. Norton & Company.

Plomp, R. (1970). Timbre as a multidimensional attribute of complex tones. In R. Plomp & G. F. Smoorenburg (Eds.), *Frequency analysis and periodicity detection in hearing* (pp. 397–410). Sijthoff.

Prajapati, D. J., Garg, S., & Chauhan, N. (2017). Interesting association rule mining with consistent and inconsistent rule detection from big sales data in distributed environment. *Future Computing and Informatics Journal*, *2*(1), 19–30. https://doi.org/10.1016/j.fcij.2017.04.003

Read, G. (1979). *Style and orchestration.* Schirmer Books.

Rimsky-Korsakov, N. (1912). *Principles of orchestration: With musical examples drawn from his own works* (M. Steinberg, Ed.; E. Agate, Trans.). Edwin F. Kalmus.

Russell, A., Goodchild, M., & McAdams, S. (2016). The Orchestration Analysis and Research Database (OrchARD). Music Perception & Cognition Lab (MPCL), McGill University.

Sandell, G. J. (1995). Roles for spectral centroid and other factors in determining blended instrument pairings in orchestration. *Music Perception*, *13*(2), 209–246. https://doi.org/10.2307/40285694

Siedenburg, K., Saitis, C., & McAdams, S. (2019). The present, past, and future of timbre research. In K. Siedenburg, C. Saitis, S. McAdams, A. N. Popper, & R. R. Fay (Eds.), *Timbre: Acoustics, perception, and cognition* (pp. 1–19). Springer International Publishing. https://doi.org/10.1007/978-3-030-14832-4_1

Straus, J. N. (2022). *The art of post-tonal analysis: Thirty-Three graphic music analyses.* Oxford University Press. http://dx.doi.org/10.1093/oso/9780197543979.001.0001

Tardieu, D., & McAdams, S. (2012). Perception of dyads of impulsive and sustained instrument sounds. *Music Perception*, *30*(2), 117–128. https://doi.org/10.1525/mp.2012.30.2.117

Temperley, D., & VanHandel, L. (2013). Introduction to the special issues on corpus methods. *Music Perception: An Interdisciplinary Journal*, *31*(1), 1–3. https://doi.org/10.1525/mp.2013.31.1.1

Touizrar, M., & McAdams, S. (2019). Aspects perceptifs de l'orchestration dans The Angel of Death de Roger Reynolds: timbre et groupement auditif. In P. Lalitte (Ed.), *Musique et cognition: Perspectives pour l'analyse et la performance musicales* (pp. 55–78). Éditions universitaires de Dijon.

Traag, V. A., Waltman, L., & van Eck, N. J. (2019). From Louvain to Leiden: Guaranteeing well-connected communities. *Scientific Reports*, *9*(1), 5233. https://doi.org/10.1038/s41598-019-41695-z

Uzzi, B., & Spiro, J. (2005). Collaboration and creativity: The small world problem. *American Journal of Sociology*, *111*(2), 447–504. https://doi.org/10.1086/432782

Zhang, S., Wu, X., Zhang, C., & Lu, J. (2008). Computing the minimum-support for mining frequent patterns. *Knowledge and Information Systems*, *15*(2), 233–257. https://doi.org/10.1007/s10115-007-0081-7