FACULDADE DE ENGENHARIA DA UNIVERSIDADE DO PORTO

A Phylogenetic approach to Folk Music Visualization

Pedro Magalhães

U.PORTO
FEUP FACULDADE DE ENGENHARIA
UNIVERSIDADE DO PORTO

Masters in Multimedia

Supervisor: Dr. Gilberto Bernardes

Second Supervisor: Hilda Romero-Velo

July 15, 2025

A Phylogenetic approach to Folk Music Visualization

Pedro Magalhães

Masters in Multimedia

Approved by:

President: Professor Doutor Jorge Manuel Gomes Barbosa

Referee: Professor Doutor Javier Félix Merchán Sánchez-Jara

Supervisor: Professor Doutor Gilberto Bernardes de Almeida

July 15, 2025

Resumo

As tradições de música folk, como as da Irlanda e da Galiza, constituem artefactos culturais complexos e em constante evolução, cuja preservação tem sido tradicionalmente assegurada através da transmissão oral e que, atualmente, são alvo de um renovado interesse na investigação digital. Contudo, a crescente disponibilidade de grandes conjuntos de dados heterogéneos coloca desafios significativos aos musicólogos, dificultando a análise e comparação de padrões estruturais entre tradições, géneros e peças individuais. Esta dissertação aborda estas limitações propondo uma ferramenta inovadora de visualização, baseada na web, que aplica metodologias de árvores filogenéticas, originalmente desenvolvidas na biologia evolutiva, à análise computacional da música folk. O sistema permite aos utilizadores explorar relações hierárquicas entre músicas, géneros e tradições, com base em características melódicas e rítmicas, possibilitando a navegação desde uma perspetiva global ao nível da tradição até à comparação detalhada de partituras. Desenvolvida através de um processo iterativo de co-design com especialistas em etnomusicologia, a ferramenta integra filtros avançados, navegação interativa e visualização lado-a-lado de partituras, facilitando a identificação de semelhanças e trajetórias evolutivas dentro e entre tradições. A avaliação, realizada com especialistas da área através de métricas de usabilidade padronizadas e feedback qualitativo, demonstra que a abordagem proposta aumenta tanto a acessibilidade como a profundidade analítica de grandes corpora de música folk. Os resultados indicam que a visualização filogenética não só complementa a análise musicológica tradicional ao revelar relações estruturais que poderiam permanecer ocultas, como também oferece uma solução escalável para futura expansão a outros conjuntos de dados e contextos de investigação. Este trabalho contribui para a etnomusicologia computacional ao unir a análise orientada por dados a uma interação intuitiva, promovendo a preservação, o estudo e a comparação intercultural do património musical intangível.

Palavras-Chave: Árvores Filogenéticas, Música Folclórica, Representações Músicais.

Abstract

Folk music traditions, such as those of Ireland and Galicia, represent complex, evolving cultural artifacts that have historically been preserved through oral transmission and are now the focus of renewed digital research efforts. However, the growing availability of large, heterogeneous datasets poses significant challenges for musicologists seeking to analyze and compare structural patterns across traditions, genres, and individual pieces. This dissertation addresses these limitations by proposing a novel, web-based visualization tool that applies phylogenetic tree methodologies, originally developed in evolutionary biology, to the computational analysis of folk music. The system enables users to explore hierarchical relationships between musical works based on melodic and rhythmic features, supporting navigation from broad tradition-level overviews to detailed score comparisons. Developed through an iterative co-design process with ethnomusicology experts, the tool integrates advanced filtering, interactive navigation, and side-by-side score visualization to facilitate the identification of pattern similarities and evolutionary pathways within and across traditions. The evaluation, conducted with domain specialists using standardized usability metrics and qualitative feedback, demonstrates that the approach enhances both the accessibility and analytical depth of large folk music corpora. Results indicate that phylogenetic visualization not only complements traditional musicological analysis by revealing structural relationships, but also offers a scalable solution for future expansion to additional datasets and research contexts. This work contributes to computational ethnomusicology by bridging data-driven analysis and intuitive interaction, supporting the preservation, study, and cross-cultural comparison of intangible musical heritage.

Keywords: Phylogenetic trees, Folk Music, Music Visualization.

Acknowledgements

First and foremost, I would like to thank my supervisor, Dr. Gilberto Bernardes, and also PhD student, Hilda Romero-Velo, for their willingness to advise and to support me and the project throughout its development. Without their constant involvement, this project would not be as complete as it ended up being.

Secondly I would like to thank all of my master's colleagues, some of which have become really close friends throughout this two year journey. Mário, João, Nawa, Shayne, thank you so much for being part of my "weekly distraction" from the making of this document, and thank you for putting up with some of my difficulties and hardships. Lastly, I have to thank all my immediate family, and all of my friends for also putting up with me through all of this process. To my parents, António and Cláudia, my sister, Clara. My friends, Rafaela, José, Luís, Miguel, Diogo, Celina, thank you all for being part of this journey, and for enduring and advising against all of my quitting attempts.

Pedro Magalhães

"One must imagine Sisyphus happy"

Albert Camus

Contents

List of Figures

List of Tables

Abbreviations

CMN	Common Music Notation
DAW	Digital Audio Workstation
ITMA	Irish Traditional Music Archive
SSM	Self-Similarity Matrix
UMAP	Uniform Manifold Approximation and Projection
t-SNE	t-Distributed Stochastic Neighbor Embedding
WWW	World Wide Web
MEI	Music Encoding Initiative
KCPS	Kilo-cycles per second
API	Application Programming Interface
IOI	Inter-Onset Intervals
JSON	JavaScript Object Notation
DOM	Document Object Model
UI	User Interface

Chapter 1

Introduction

This dissertation investigates the computational analysis and comparison of folk music scores, aiming to uncover structural patterns through less common visualization approaches. **Section** 1.1 introduces the motivation and context, highlighting the need to preserve and analyze folk music as a dynamic cultural artifact. **Section** 1.2 outlines the research objectives, focusing on developing methods to compare multiple music scores and visualize their similarities and differences. **Section** 1.3 provides an overview of folk music, emphasizing its cultural significance and the challenges of representing its diversity in a computational framework. **Section** 1.4 describes the dataset of folk music scores used in this project, including its sources, features, and preprocessing steps. Finally, **Section** 1.5 presents the structure of the dissertation, guiding the reader through the research process and its contributions to musicology and data-driven analysis.

1.1 Context and Motivation

This dissertation is situated within the broader framework of the **EA-DIGIFOLK** project, an Ibero-American initiative focused on the digital collection, analysis, and dissemination of folk music. The project aims to create a publicly accessible platform that curates and presents information on European and Latin American folk music, facilitating musical and ethnomusicological analysis for both academic and community use. Within this context, the Sound and Music Computing Group at FEUP and INESC TEC has been tasked with developing algorithms to identify similarities across folk songs, using both melodic and harmonic features.

The project utilizes datasets from **I-FOLK**, **"(A Música Portuguesa a Gostar Dela Própria)"**, the **Irish Traditional Music Archive** (ITMA), the **Folkoteca Galega**, and others. These comprise a corpus of over 400 songs (and expanding) from Portuguese, Italian, and Spanish folk traditions. These datasets incorporate detailed musical analyses of mode, rhythm, and phrasing, providing a rich foundation for a visualization that aims to reveal structural patterns, interconnections across traditions, among others.

My involvement in this project is driven by its potential to create visual representations that enhance the workflow of researchers and musicologists. It is important to emphasize that the goal is not to automate or replace their work but to provide tools that enable clearer and more intuitive navigation of large datasets. By offering new ways to explore and interpret folk music, this project seeks to support and enrich the research processes of scholars and experts in the field.

1.2 Objectives and Problem

As previously mentioned, EA-DIGIFOLK comprises datasets from several different traditions, resulting in a vast and continually expanding corpus of music. While this abundance of material is expected in a project of such broadness, encompassing diverse traditions and recovered data, it also presents a significant challenge for researchers. Without effective tools to navigate and interpret it, the sheer scale of the dataset can become a serious obstacle to meaningful analysis.

As the project progresses and more data is collected, this work becomes increasingly relevant. The investigation being conducted addresses the challenges posed by large datasets by creating visual representations of a musical corpus. This approach enables the analysis of vast quantities of data, by making it easier to detect patterns, and identify features that stand out. From this perspective, the research that will be conducted is essential not only for the development of the project, but also for the valuable insights it will generate. The main objective of this project is to create a visualization to ultimately organize all the information available within the project's scope.

By doing so, we aim to bridge the gap between traditional music preservation methods and modern computational tools. Additionally, this work seeks to address the problem of scalability when dealing with large datasets, ensuring that the proposed solutions are not only fit for the current data, but built for future expandability.

Through this dissertation, we will evaluate existing visualization methods, identify their limitations, and propose new approaches tailored to the unique characteristics of folk music. Ultimately, we aim to contribute to the preservation and understanding of folk music while also applying visualizations that are not commonly used for folk music, into it.

1.3 Overview of Folk Music and its Cultural Importance

Bruno Nettl is a central figure in this chapter, given that he advanced the pioneering work of **George Herzog** in ethnomusicology. Building on Herzog's foundational contributions, Nettl refined and expanded the field, making his work essential to this study (Nettl, 2025).

As Nettl defined **folk music** as a type of traditional and generally rural genre of music that historically lives solely off of oral tradition. This way, it was passed down generationally or through small social groups (village communities). Having these characteristics, folk music inherently introduces a great deal of change, regardless of how far it roots. Being it for people on these communities forgetting parts of it, for them wanting to improve it, or to make it sound more similar to another song they were listening to at the time. Folk and primitive music, then, have for us the fascinating quality of being both old and contemporary, of being representative of a people's ancient traditions as well as an indicator of their current tastes (Nettl, 1965).

The preservation and study of folk music have long been a focus of ethnomusicologists. At the beginning of the 20th century, early European ethnomusicologists were driven by a pedagogical interest in collecting and documenting traditional music. Among these figures, Zoltán Kodály, a renowned musician and pedagogue, stands out. Kodály not only contributed to the collection of folk music but also developed an influential method to simplify and enhance musical education. His work exemplifies the dual goals of preserving cultural heritage and making music accessible to broader audiences, a theme that resonates with the objectives of the EA-DIGIFOLK project (Göktürk, 2012).

Early folk music preservation was rooted in ethnographic fieldwork, with figures like Béla Bartók using wax cylinders and Alan Lomax employing instantaneous disc recordings to document performances in-site. These technologies enabled the capture of not only melodies but also the subtle performative and contextual elements essential for understanding oral transmission and cultural context (kolovos, 2010). By the mid-20th century, institutions such as the Library of Congress' Archive of Folk Culture established systematic metadata practices, recording performer demographics, geographic origins, and social context, which greatly improved the analytical value and comparability of collections (kolovos, 2010).

Technological advances have since transformed preservation. Projects like the Silesian Digital Library digitized manuscripts, unpublished texts, and reel-to-reel tapes, using modern notation and audio restoration to reconstruct lost details (Krajewska, 2020). Meanwhile, the Hua'er folk song database in China applies machine learning to classify regional subtypes and trace stylistic diffusion, exemplifying how computational tools address conservation challenges and enhance cross-cultural research (Guangchao et al., 2024).

In recent years, a diverse array of international initiatives has transformed the preservation and promotion of folk music heritage. These efforts range from large-scale digitization of historical recordings and manuscripts to the creation of sophisticated online platforms that ensure both public accessibility and robust scholarly research. By consolidating and standardizing data from such varied sources, projects like EA-DIGIFOLK not only safeguard the musical traditions of regions such as Galicia and Ireland but also enable comprehensive analysis and cross-cultural comparison. This integrated approach is essential for advancing our understanding of folk music's complexity and ensuring its vitality for future generations.

1.4 Datasets

The primary objective of this project is to enhance the visual accessibility of existing and future datasets from various folk traditions, thereby making them more usable for scholars and researchers in the field. Within the scope of this dissertation, the initial focus will be placed on a subset of the broader EA-DIGIFOLK project: datasets from the Galician and Irish traditions. Together, these two datasets comprise 600 pieces, including both unique songs and variations of the same song across traditions. These works span 21 distinct genres, offering a rich and diverse foundation for the development and testing of visualization strategies.

A key reason for selecting the Galician and Irish datasets is their shared historical and cultural roots as part of the broader Celtic tradition. Both regions have musical lineages that trace back to ancient Celtic cultures, resulting in notable similarities in melodic structures, rhythmic patterns, and even certain instruments. This common heritage makes the comparison between Galician and Irish folk music particularly relevant, as it enables the study of how a shared origin has evolved into distinct yet related musical expressions. By focusing on these two traditions, the project not only leverages a rich dataset but also addresses important questions about musical evolution, cultural exchange, and the preservation of intangible heritage.

Although this dissertation focuses on the Galician and Irish datasets, it is important to note that the broader EA-DIGIFOLK project also encompasses datasets from Portuguese and selected Latin American folk traditions. The integration of these additional datasets, along with others, will be discussed later in this document, as part of the project's ongoing expansion.

1.5 Dissertation structure

This dissertation is organized into six main chapters, each addressing a distinct stage of the research process and collectively guiding the reader from the theoretical foundations to the practical development of visualization strategies for musicological analysis.

Chapter 1: Introduction

The opening chapter establishes the context and motivation for the research, situating it within the EA-DIGIFOLK project and outlining the challenges and objectives associated with the digital collection of folk music. It provides an overview of the cultural significance of folk music, introduces the datasets used primarily from Galician and Irish traditions, and defines the research questions that guide the subsequent work.

Chapter 2: Review of Music Visualization Techniques

This chapter presents a comprehensive review of existing music visualization methods relevant to the analysis of folk music. It begins with a scoping review of the literature, detailing the methodology used to identify and categorize key visualization strategies. The chapter then explores the structured elements of music, distinguishing between symbolic and audio representations, and discusses a range of visualization techniques, including those for individual melodies, complete songs, and large corpora. Special attention is given to the adaptation and evaluation of these techniques for the unique characteristics of folk music datasets.

Chapter 3: Prototyping and Co-Design Process

The third chapter documents the practical development of the visualization tool. It describes the iterative process of creating initial mockups, the rationale behind design decisions, and the collaborative approach taken in partnership with supervisors and domain experts. The chapter details the structure and outcomes of a co-design session, highlighting how expert feedback was integrated to refine the prototypes and ensure their relevance and usability for musicological research.

Chapter 4: Prototype Implementation

Chapter four translates design concepts into a functional visualization system. It begins with an architectural overview that establishes the technical foundation for representing phylogenetic relationships in folk music. The chapter examines the data preprocessing pipeline, detailing how musical scores are analyzed and transformed into structures suitable for visualization. It then explores the implementation of key components, from interactive tree rendering to score comparison features, highlighting technical challenges encountered and solutions developed. The chapter concludes with a post-implementation evaluation, documenting expert feedback on the deployed system and explaining how this informed iterative refinements to enhance the tool's utility for musicological research.

Chapter 5: Evaluation

The fifth chapter presents a comprehensive assessment of the implemented visualization tool. It begins by examining the technical challenges encountered during development and the strategies employed to overcome them. The chapter then analyzes feedback gathered from domain experts during follow-up evaluations, comparing user experiences with the implemented system against initial expectations from the co-design sessions.

Chapter 6: Conclusion

The sixth and final chapter presents the conclusion of this dissertation. It outlines potential directions for future work, discusses features that could not be implemented within the project's timeframe, and proposes enhancements to further expand the tool's capabilities for folk music analysis and comparison. Additionally, this chapter evaluates the extent to which the project's core objectives were achieved, highlighting both the strengths and limitations of the current implementation.

Chapter 2

Review of Music Visualization Techniques

This chapter provides a comprehensive review of music visualization techniques relevant to the analysis of folk music, establishing a foundation for the development of new visualization strategies tailored to this genre. The review is structured into four main sections. First, **Section** 2.1 outlines the methodology used to identify and categorize existing visualization approaches. **Section** 2.2 explores the structured elements of music, clarifying the conceptual framework that underpins the techniques discussed later. **Section** 2.3 then examines a range of music visualization methods, distinguishing between symbolic and audio representations. Finally, **Section** 2.4 analyzes these techniques and identifies opportunities for developing new visualizations specifically suited to folk music, setting the stage for the practical work described in subsequent chapters.

2.1 Scoping Review

To establish a strong theoretical and methodological foundation for this project, a scoping review was conducted. A scoping review is a type of literature review that aims to map the key concepts, types of evidence, and gaps in a given research area by searching, selecting, and synthesizing existing knowledge. Unlike systematic reviews, which often seek to answer narrowly defined research questions, scoping reviews are particularly useful in fields that are complex or have not been comprehensively reviewed before. Given the interdisciplinary nature of this project, spanning folk music studies, music visualization, and computational music analysis, a scoping review was deemed the most suitable approach.

This methodology allowed for an inclusive and exploratory examination of a broad field, helping to identify major trends, theoretical frameworks, and methodological tools that could inform both the conceptual and practical dimensions of the dissertation.

2.1.1 Review Strategy and Methodology

The review process began with a broad search for academic papers and articles related to folk music, music visualization, and the intersection of the two. Google Scholar was used as the primary search engine due to its extensive indexing of academic sources. Keywords such as "folk music", "music visualization techniques," and "music representations" were used to identify relevant literature. Results were filtered by publication date to prioritize recent contributions (primarily from the last decade), though foundational older works were also considered when relevant.

To strengthen the scope and depth of the review, a "snowballing" technique was also employed, examining the reference lists of key state-of-the-art papers to discover additional important sources. This process ensured the inclusion of both contemporary and foundational works that may not have appeared in the initial keyword-based searches.

In total, more than fifty sources were selected for analysis, although only about thirty five were cited within this dissertation work. These were primarily composed of peer-reviewed papers and scholarly articles, and were organized in a structured spreadsheet according to publication year, category, and thematic focus. The literature was grouped into three main categories, namely, Folk music studies, Interdisciplinary works relevant to multiple aspects of the project, and Music visualization research. Given the diversity of music visualization approaches, this final category was further divided into two subcategories: visualizations focused on small-scale musical structures (e.g., melody, phrases, full songs), and visualizations of large-scale corpora, often employing clustering or dimensionality reduction techniques.

By combining systematic searching, citation tracking, and structured categorization, the scoping review provided a well-rounded overview of the literature. This process ensured that the current research is grounded in the most relevant theoretical and technical knowledge available, and helped identify the tools and strategies most appropriate for visualizing structural relationships within folk music corpora.

2.2 Music Information Retrieval

As the name implies, **Music Information Retrieval** (hereafter, **MIR**), is a field dedicated to developing computational methods for analyzing, organizing, and retrieving musical information. Music, as an art form, is inherently structured, with its complexity arising from the interplay of elements such as rhythm, harmony, and timbre. To analyze and understand music computationally, it is essential to break it down into its fundamental components, which can be categorized into the aforementioned, structured elements. These elements provide a framework for examining music at different levels of abstraction, from objective acoustic properties to high-level perceptual and semantic features that align with human interpretation (Lerch, 2023).

In this section, we explore structured elements through the lens of **low-level** and **high-level audio descriptors**, drawing on insights from MIR research. Low-level descriptors capture the immediate, physical properties of the audio signal, such as spectral shape, energy distribution, and temporal dynamics. These features serve as the building blocks for more sophisticated analyses. In contrast, high-level descriptors abstract these low-level features into musically meaningful concepts, such as tempo, key, and genre, which align more closely with human perception and interpretation (Müller, 2021).

The distinction between low, mid, and high-level musical descriptors reflects the hierarchical nature of music itself, where each level captures different layers of musical information. **Low-level** features describe the fundamental properties of individual sounds, and the physical attributes of the wave, such as the Mel-frequency Cepstrum Coefcient (MFCC) and Zero Cross Rate (ZCR) (Corrêa and Rodrigues, 2016). These features provide detailed, raw data about the acoustic or symbolic elements of music but do not convey relationships between sounds or broader musical context. For example, MFCC is a widely used feature extraction technique that represents the short-term power spectrum of a sound by applying a discrete cosine transform to the logarithm of the power spectrum, which has first been mapped onto the nonlinear mel frequency scale. This process effectively models the way the human auditory system perceives sound, emphasizing perceptually relevant frequency components (Ali et al., 2021).

High-level features capture the overall organization, expressive qualities, and stylistic characteristics of a musical piece. These features operate at a broader, more abstract level than low-level descriptors, reflecting aspects such as melody, harmony, rhythm, and emotional content. Examples of high-level features include key changes, chord progressions, rhythmic patterns, and dynamic contours (Corrêa and Rodrigues, 2016). Unlike low-level features, which focus on the physical properties of individual sounds, high-level features describe relationships and structures that emerge across larger segments of music. For instance, key changes in harmony can indicate shifts in mood or tension, while patterns in rhythm and dynamics contribute to the perceived style and emotional impact of a piece. These features are essential for understanding the musical narrative and for tasks such as genre classification, mood detection, and comparative music analysis, as they align closely with human perception and interpretation of music.

Mid-level features go beyond isolated sounds by representing relationships and patterns within the music. Examples include the beats-per-minute (BPM), pitch histogram (PH) and chord sequences (CS) (Corrêa and Rodrigues, 2016). These descriptors capture local structures and recurring motifs, such as a chord progression or a rhythmic groove, which contribute to the music's internal coherence and style.

By leveraging descriptors across these three levels, we can enhance music categorization and develop more effective visualizations that connect raw audio or symbolic data to meaningful musical insights. Low-level features offer the detailed raw material, mid-level features reveal patterns and relationships, and high-level features provide interpretive context. This multi-layered approach bridges the gap between technical analysis and the rich, expressive experience of music (Corrêa and Rodrigues, 2016).

This section also delves into **music visualization techniques**, which can be divided into two major categories: **Symbolic Representations** and **Audio Representations**. Symbolic representations provide a direct depiction of musical elements, such as notes, rhythms, and harmonies, often aligning with traditional musical notation. These representations are highly intuitive for end users, offering a clear and structured view of the music. In contrast, audio representations deal with the raw audio signal, which is inherently more abstract and complex. These representations require sophisticated techniques to extract and convey information about timbre, dynamics, and spectral content in a meaningful way (Lerch, 2023). The differences in their characteristics necessitate distinct visualization approaches, which will be explored in detail throughout this document.

2.3 Music Visualization Techniques

Music visualization techniques hold significant historical importance, as they laid the foundation for the development of sheet music, one of the most enduring and universal systems for representing music. These techniques emerged as a means to bridge the gap between auditory experiences and visual representation, enabling the preservation, communication, and analysis of musical ideas (Dean, 2009).

Guido d'Arezzo, an 11th-century Italian monk, is widely regarded as a pioneer in the field of music visualization. Guido is credited with creating one of the earliest known systems for visually representing melodies, effectively connecting musical sounds to written symbols. His innovations, such as the Guidonian hand and the development of musical notation, revolutionized the way music was taught, learned, and transmitted, marking a critical step in the evolution of music visualization (Miller, 1973).

However, the attribution of the invention of the modern music staff to Guido still remains a subject of debate and continues to be widely discussed in the scientific community, to this day (Miller, 1973).

> "CMN reflects many centuries of accumulated user feedback and collective wisdom synthesized in a flexible and adaptive representation" (Isaacson, 2005)

The origins of Common Music Notation (most commonly known as musical scores) can be traced back to the Middle Ages, with significant refinements occurring during the Renaissance and incremental modifications continuing to the present day. Musical Scores emerged as a solution to preserve the extensive *plainchant repertoire* of the Roman Catholic Church, which had previously been transmitted through **oral tradition**. Its primary purpose was, and remains, to facilitate the performance of musical compositions by providing performers with a clear and structured guide (Figure 2.1 serves as an example of one of the most well-known music pieces) (Miller, 1973; Isaacson, 2005; Dean, 2009).

Much like the written word enabled the codification of languages and the emergence of literature, the development of music notation allowed for the conceptualization of a musical work as a reproducible artifact that could be reliably passed down through generations. This shift not only preserved musical ideas but also enabled the creation of a shared cultural and artistic heritage, laying the groundwork for the evolution of Western music (Dean, 2009).

Figure 2.1: Beethoven's Für Elise, a well-known example of a musical score

While music scores represented a significant advancement in music preservation, its development and use were historically restricted to clerical and elite circles, particularly during the Middle Ages. This "gatekeeping" of musical knowledge meant that musical scores did not extend to the broader folk *strata*, leaving much of the music created and performed by ordinary people unrecorded and reliant on oral tradition. As a result, although music scores facilitated the preservation and transmission of sacred and elite musical works, it did not encompass the full diversity of musical expression across all social and cultural contexts (Feld and Nettl, 1986).

In 1967, the first computer generated music visualization tool was showcased in "Fritz Winckel (1967) Music, Sound and Sensation: A Modern Exposition" (Figure 2.2). This was a physical 3D spectrogram showing a frequency analysis of an 8-second recording of Beethoven's Eighth Symphony. Winckel's book didn't provide any information as per the visualization's origins, or how it was made (Dragicevic, 2018).

Nowadays, music visualizations have become more ubiquitous, proving once more their usefulness by providing the user with intuitive ways to analyze, edit, and interact with music. Be it from the piano roll present in almost all of the DAWs (Digital Audio Workstation), to waveforms in sound recording software (Cooper et al., 2006; Nieto et al., 2020).

In this project, we focus on developing analytical and objective approaches to visualizing folk music. While detailed documentation on folk music visualization is limited, it is essential to evaluate all available materials for their adaptability to this genre. We will explore multiple visualization methods tailored to each abstraction type, ensuring they align with the structural and expressive qualities of folk music (Panteli et al., 2018).

2.3.1 Symbolic Representations

Symbolic representation refers to the use of visual symbols to encode musical information, such as pitch, rhythm, and dynamics, in a structured and interpretable format. The most widely recognized form of symbolic representation is music scores, which has been the standard for Western music for centuries. Music scores provide a precise and standardized way to represent musical ideas, enabling composers, performers, and scholars to communicate and preserve musical works across time and space (Isaacson, 2005).

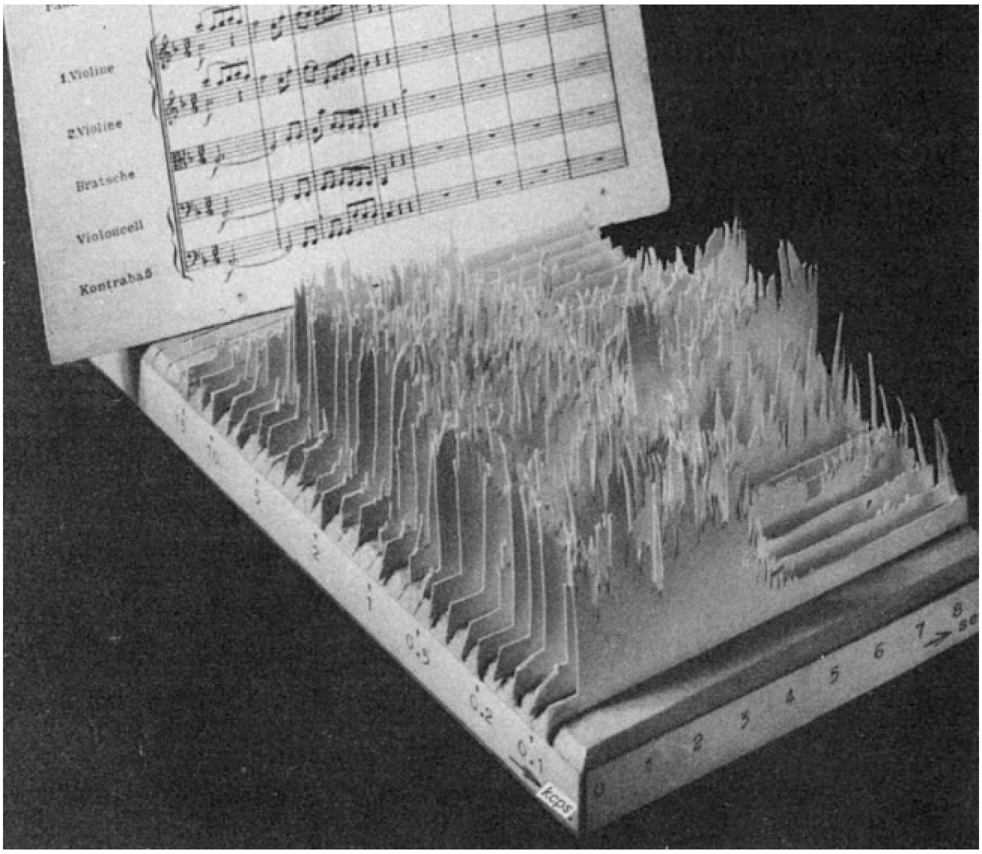

Figure 2.2: Sonagram, beginning of Beethoven's Eight Symphony, displaying the overtone structure up to 12 ksps (Kilo-cycles per second). Taken from (Dragicevic, 2018).

Symbolic representations are particularly valuable for tasks that require precise control over musical elements, such as score-based analysis and music transcription. They also serve as a foundation for MIR systems, providing unparalleled precision for tasks requiring granular control over musical elements, enabling the extraction of quantitatively exact parameters that form the foundation of computational analysis. These include, absolute pitches, intervalic relationships, inter-onset intervals (IOI), note durations and metric positions (Liu et al., 2021). Although they contain all of this data, they still struggle to capture the nuances of timbre, expression, and improvisation, which are better represented through audio signals or hybrid approaches (Liu et al., 2021).

In summary, symbolic representations play a crucial role in music visualization by providing a structured and precise framework for encoding and analyzing musical information. Notably allowing for highly controlled evaluation of similarity between pieces, making them particularly valuable for comparative analysis. In the following sections, we will examine how different abstractions derived from folk music sheet music transcriptions within a dataset can be represented and analyzed using this approach.

2.3.1.1 Melody and Phrase Elements

The use of **augmented scores** for visualizing melody and phrase elements in sheet music is grounded in their ability to enhance traditional notation with additional layers of analytical and expressive detail (Figure 2.3) (Cantareira et al., 2016; Khulusi et al., 2020).

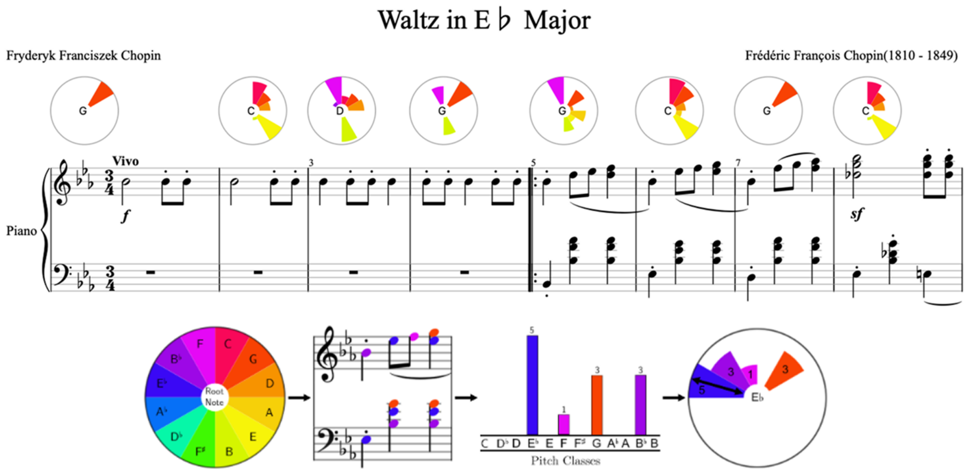

Figure 2.3: Nightingale rose of the circle of fifths, by Miller et al. (Screenshot). Each musical compass has a Nightingale-rose above, summarizing the notes. A better explanation to how the tool's visualization language works is below. Tool available at: https://musicvis.dbvis.de/app/fingerprint/ last accessed on: 19/06/2025 at 16:36 GMT.

Much like the layered and multidimensional approaches described in visualization research, augmented scores provide a framework where pitch, rhythm, dynamics, and phrasing can be intuitively represented through visual means. In this context, color gradients serve as tools to highlight melodic contours, expressive nuances, and structural relationships. This visual approach allows for a more immediate and comprehensive understanding of the music's organization, enabling scholars to identify patterns, variations, and connections that might otherwise remain hidden in conventional notation. The ability to 'see' the interplay of melodic and phrase elements fosters deeper engagement. By integrating modern visualization techniques with the precision of traditional notation, augmented scores offer a powerful tool for exploring the intricacies of musical structures in a way that is both accessible and insightful (Lima et al., 2022).

2.3.1.2 Complete Songs

The use of arc diagrams for visualizing complete sheet music is grounded in their ability to effectively represent structural and relational features of musical compositions (Figure 2.4) (Wattenberg, 2002).

Arc diagrams provide a spatial framework where relationships between musical elements can be intuitively visualized. In this case, arcs represent connections or interactions between notes, phrases, or sections, with their positioning and curvature conveying information about consonance, dissonance, and structural coherence (Cooper et al., 2006). This visual approach allows for a quick

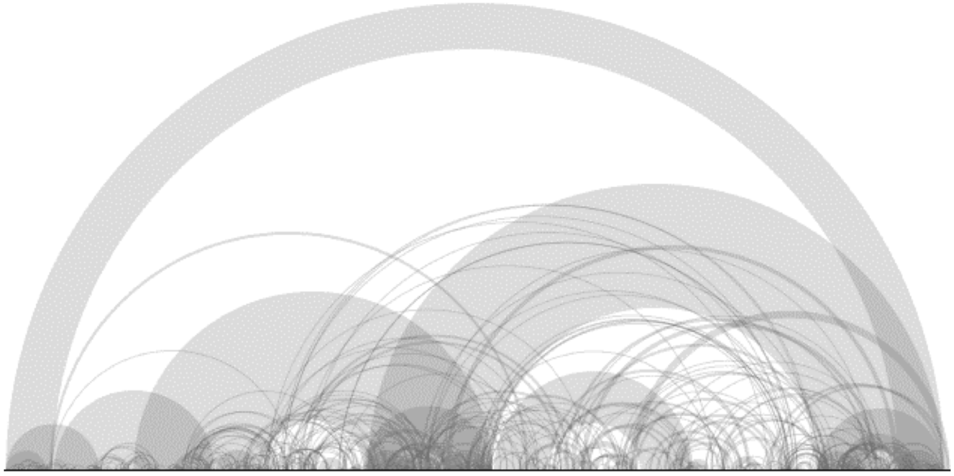

Figure 2.4: Arc Diagram of Beethoven's Für Elise. Taken from (Wattenberg, 2002)

yet comprehensive overview of the musical structure, enabling scholars to identify patterns, transitions, and relationships that might otherwise remain obscured in traditional notation. Furthermore, the clarity of arc diagrams supports both the creative and analytical processes—whether composing new music or editing existing works, the ability to see the structure of a piece facilitates informed decision-making and fosters a deeper understanding of the musical material (Isaacson, 2005). By translating complex musical relationships into an accessible visual format, arc diagrams serve as a powerful tool for engaging with the intricacies of complete musical works.

2.3.2 Audio Representations

Waveform representation refers to the visualization of audio signals as a function of amplitude over time, providing a direct and, somewhat, intuitive depiction of sound. Unlike symbolic representations, which encode musical information through abstract symbols, waveforms capture the raw acoustic properties of audio, including dynamics, timbre, between others. Although this makes waveforms particularly useful for analyzing and editing audio signals in their most fundamental form, it also presents new challenges in the visualization realm (Müller, 2021; Lerch, 2023).

Waveforms are widely used in sound recording and production software, where they serve as a primary tool for visualizing and manipulating audio (Figure 2.5). They allow users to identify key features such as volume changes and silences. Additionally, waveforms are often used in Music Information Retrieval (MIR) systems to extract low-level features such as energy, zero-crossing rate, and spectral content, which can be further processed to derive higher-level musical information (Blaszke and Koszewski, 2020; Lerch, 2023).

However, waveforms have limitations when it comes to representing abstract musical concepts such as pitch, harmony, and rhythm. While they excel at capturing the physical characteristics of sound, they lack the structured and interpretable nature of symbolic representations. For example, a waveform cannot directly convey the melodic or harmonic structure of a piece of music, making it less suitable for tasks like score-based analysis or transcription (Müller, 2021; Lerch, 2023).

Figure 2.5: Waveform of Beethoven's Für Elise. Taken from (int).

In summary, waveform representation plays a vital role in music visualization, offering a direct and detailed view of audio signals. While they are less suited for encoding abstract musical concepts, they provide invaluable insights into the acoustic properties of sound. In the following sections, we will explore representations that can arise from several abstractions of folk music field recordings, within a dataset (Cooper et al., 2006).

2.3.2.1 Melody and Phrase Elements

Melody and phrase elements represent the lowest levels of abstraction in this project and, as such, will feature the most direct forms of visualization. For this waveform representation segment, the use of color to encode different types of descriptors is particularly relevant. Colored waveforms, for instance, can effectively convey pitch information, with variations in hue or intensity reflecting changes in pitch height or contour. This approach not only enhances the visual appeal of the waveform but also provides an intuitive way to perceive melodic patterns and transitions at a glance (Cooper et al., 2006; Malandrino et al., 2015; Khulusi et al., 2020).

Building on this idea, other descriptors such as dynamics, timbre, or harmonic content could also be represented through color gradients or additional visual layers. For example, brighter colors might indicate louder dynamics, while shifts in saturation could reflect timbral changes. By integrating multiple descriptors into a single visualization, the waveform becomes a rich, multidimensional representation of the melody, offering both analytical depth and immediate perceptual clarity (Malandrino et al., 2015; Lima et al., 2022; Cooper et al., 2006).

2.3.2.2 Complete Songs

> "Music pieces often show some coherence or similarity over their durations, and this is an example of how visual representations can be used to visualize patterns within a music piece [...] using a matrix; similar regions are represented with vivid colors while dissimilar regions are darker." (Lima et al., 2022)

Self-similarity matrices (SSM's) provide a powerful and structured approach to visualizing complete songs through waveform representation (Figure 2.6). The underlying principle involves transforming the music signal into a sequence of meaningful features and then comparing each element of this sequence with all other elements. This process generates a self-similarity matrix, which effectively captures the internal relationships and patterns within a musical piece. By mapping these relationships, SSM's reveal essential structural features such as repetitions, variations, and transitions, offering a comprehensive overview of the song's organization (Louie, 2020; Cooper et al., 2006; Foote, 1999; Müller, 2021).

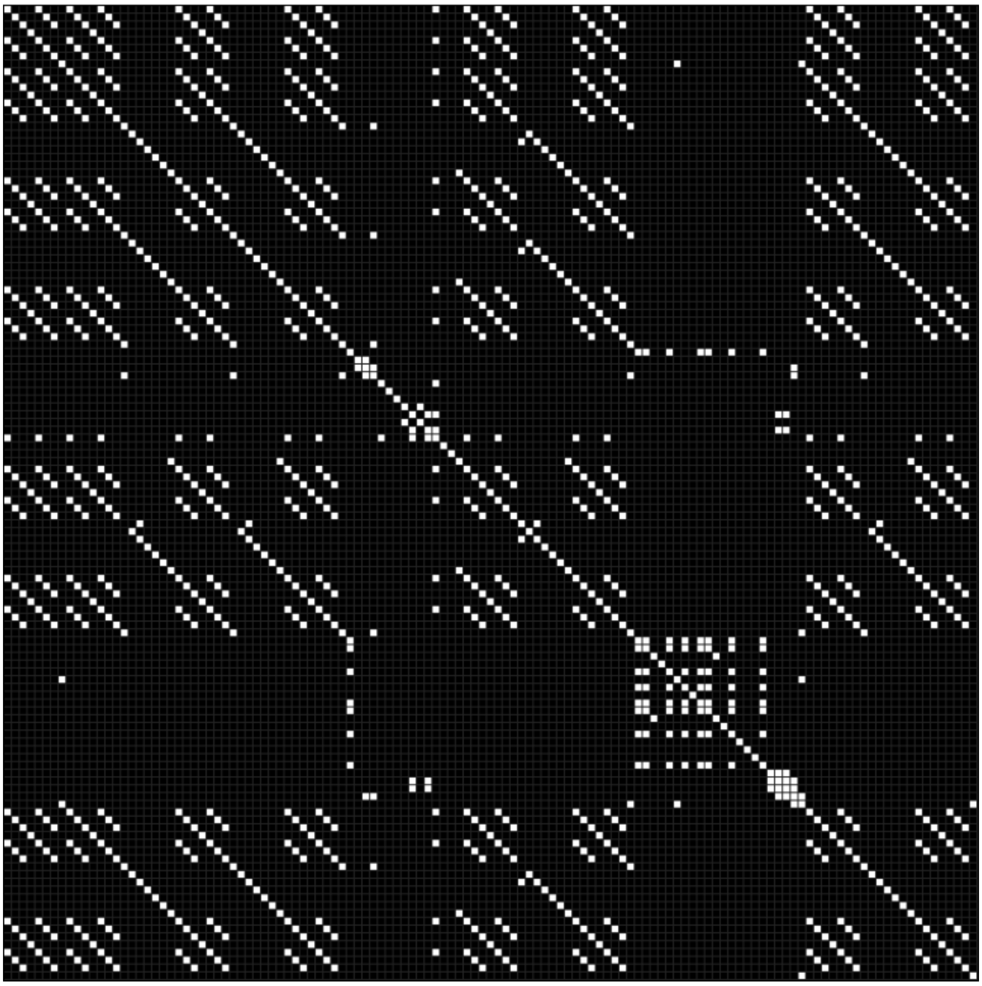

Figure 2.6: Self-Similarity Matrix of Beethoven's Für Elise. Taken from (Louie, 2020).

A key strength of SSM's lies in their ability to highlight the coherence and repetition that are often present in music. Musical pieces frequently exhibit patterns of similarity over their durations, and visual representations like SSM's make these patterns immediately apparent. In such matrices, regions of high similarity are represented with vivid colors, while non-similar regions appear darker, creating a clear and intuitive visual distinction. This approach not only aids in identifying structural elements like repeated motifs or contrasting sections but also provides a deeper understanding of the compositional logic and expressive qualities of the music (Foote, 1999; Khulusi et al., 2020; Lima et al., 2022).

In this project, SSM's serve as a versatile and objective tool for visualizing complete songs, particularly in the context of folk music, where structural patterns and repetitions are often central to the genre. By employing this method, we aim to create visualizations that not only capture the temporal and structural dynamics of a piece but also facilitate a more nuanced analysis of its underlying patterns and relationships (Panteli et al., 2018; Liu et al., 2021).

2.3.3 Corpus Visualization

Corpus visualization is not segmented in the same way as the other types of abstractions, because it makes sense to have a common visualization basis for both **Symbolic Representations**, and for **Audio Representations**.

The Music Encoding Initiative (MEI) is a community-driven, open-source effort to define a system for encoding musical documents in a machine-readable structure. Within this framework, MEI describes a **corpus** as a collection of musical data, often curated and organized with a specific purpose in mind. Unlike arbitrary assemblages of musical works, a corpus is typically designed to highlight particular features or to capture the diversity within a group of related pieces. What sets a corpus apart is its intentional design: the selection and arrangement of its components are guided by a set of deliberate criteria, ensuring that the collection serves as a coherent and meaningful resource for analysis. This thoughtful construction allows researchers to engage with the corpus in ways that align with its intended goals, whether to explore specific musical traits, trace stylistic developments, or examine variations across a body of work (Carvalho et al., 2021).

However, working with large collections poses challenges, particularly in classifying and aggregating vast amounts of data. Visualizations must go beyond simple file listings. **Clustering** is an unsupervised learning technique, and it addresses this by grouping objects into clusters where similarities within clusters are maximized, and differences between clusters are minimized. This method uncovers patterns and structures within complex datasets, making it invaluable for organizing and interpreting musical data (Irani et al., 2016; Wijaya and Oetama, 2021).

Algorithms like **UMAP** (Uniform Manifold Approximation and Projection) and **t-SNE** (t-Distributed Stochastic Neighbor Embedding) are particularly effective for non-linear dimensionality reduction. As per the example on Figure 2.7, where we can see how both of these clustering philosophies influence the final arrangement of each cluster. They project high-dimensional data into a simpler 2D plane, revealing underlying patterns and relationships that might otherwise remain hidden. This simplification is crucial for visualizing multidimensional musical data, enabling researchers to identify trends, similarities, and differences within a corpus (Carvalho et al., 2021; Khulusi et al., 2020; Lima et al., 2022). By encoding complex data into a 2D plane, these visualizations not only enhance accessibility but also support deeper analytical and creative engagement with the material.

Another way to visualize a large corpus is by using **phylogenetic trees**. These are generally applied in genetics to estimate of the relationships among taxa (or sequences) and their hypothetical common ancestors. Because of this, phylogenetic trees provide a structured and hierarchical representation of musical relationships, where each branch illustrates connections based on shared or derived traits. This method groups pieces into clusters of similarity, emphasizing common origins and divergences, while minimizing unrelated or dissimilar groupings. By doing so, phylogenetic trees uncover patterns and structures within complex datasets, making them invaluable for

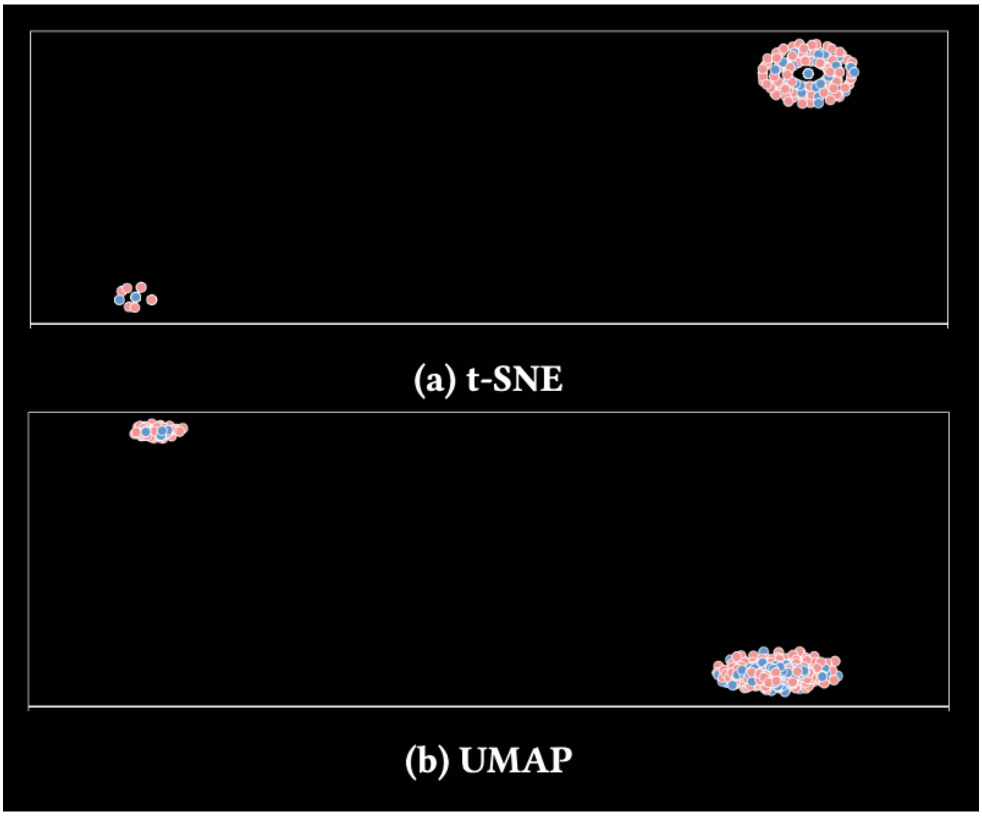

Figure 2.7: 2-D visualizations of the i-Folk database using t-SNE (a) and UMAP (b). Points closer together represent songs from the same country, while colors denote genre (red: Corro, blue: Cuna). Taken from (Carvalho et al., 2021).

analyzing the distribution of musical features across a corpus (Thul and Toussaint, 2008; Panteli et al., 2018; Liebman et al., 2012).

Phylogenetic trees are particularly effective for revealing temporal, geographic, and stylistic relationships. For example, in a corpus of folk music, a tree could trace the spread of a melody across regions, with branches showing variations in, for example, rhythm. The length and structure of the branches provide insights into the degree of similarity or divergence, offering a visual measure of how closely related pieces are. This approach not only highlights shared roots but also reveals unique adaptations that emerge in different cultural or historical contexts (Thul and Toussaint, 2008; Liebman et al., 2012).

2.4 Analysis and Opportunities for New Folk Music Visualizations

Building upon the established visualization techniques surveyed in earlier sections, spanning symbolic score representations, audio waveform analyses, and corpus clustering methods, collaborative discussions with EA-DIGIFOLK researchers highlighted phylogenetic trees as a promising yet underutilized approach for folk music analysis.

Phylogenetic trees, originally developed in evolutionary biology to model species patterns, provide hierarchical representations of divergence from common ancestors through branching structures. In computational musicology, this paradigm maps effectively to the study of melodic evolution across folk traditions, where variants emerge through oral transmission and regional adaptation (Thul and Toussaint, 2008; Panteli et al., 2018; Liebman et al., 2012). Although previously mentioned as being underused, some of the literature gives some data visualization examples using phylogenetic trees. One such example is shown in Figure 2.8, which in the context of (Liebman et al., 2012), represents performers' presumed association with different musical schools, categorized into five main groups, according to the primary teachers under whom they studied.

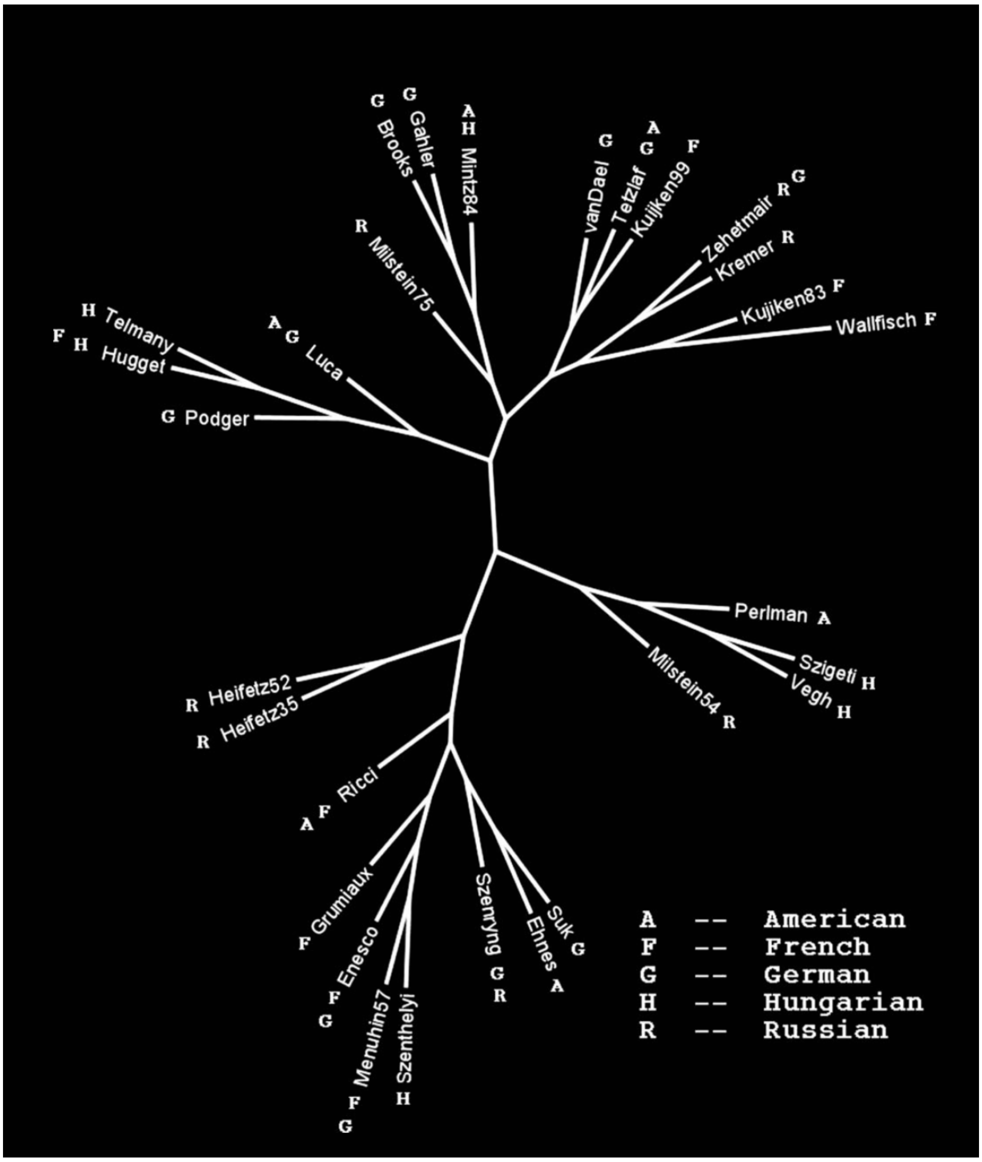

Figure 2.8: Phylogenetic Tree example. This tree depicts performers' alleged affiliation to musical schools (roughly divided into five categories) based on their primary teachers. Taken from (Liebman et al., 2012)

Theoretical Foundations and Functional Advantages

Phylogenetic trees address two critical limitations of conventional corpus visualization techniques:

1. Non-linear relationship modeling: Unlike dimensionality reduction methods (e.g., UMAP, t-SNE) that project high-dimensional data onto static planes, phylogenetic trees preserve directional divergence pathways. This allows researchers to trace hypothetical evolutionary trajectories between variants, distinguishing convergent adaptations from shared ancestral traits. For instance, a Galician "muiñeira" and an Irish "reel" might occupy proximal coordinates in a UMAP plot due to rhythmic similarities, but a phylogenetic tree could reveal independent origins through distinct branching patterns (Khulusi et al., 2020; Carvalho et al., 2021; Müller, 2021; **?**).

2. Multi-scale analytical capacity: The nested hierarchy supports simultaneous examination of macro-level tradition comparisons (e.g., Iberian vs. Celtic) and micro-level variant differentiation within genres. This aligns with folk music's inherent structural organization, where traditions bifurcate into genres, which further subdivide into regional or performative variants.

Implementation Considerations

Key implementation parameters derive from computational biology and MIR:

1. Distance metrics: Needleman-Wunsch sequence alignment scores quantify melodic dissimilarity by weighting pitch intervals, rhythmic ratios, and phrase structures. Hybrid metrics combining symbolic features (e.g., mode, ornamentation patterns) with audio descriptors (spectral centroid, onset density) enable multimodal comparisons (Cooper et al., 2006; Blaszke and Koszewski, 2020).

2. Rooting strategies: Outgroup rooting using archaeologically attested melodies (e.g., medieval cantigas) provides temporal orientation, while midpoint rooting optimizes topological balance for cross-tradition analyses (Panteli et al., 2018).

3. Branch length encoding: Proportional to dissimilarity scores, branch lengths visually weight divergence magnitude. In Galician mazurcas, elongated branches often correlate with rhythmic ornamentation shifts, while shorter branches indicate conserved phrase structures (Liebman et al., 2012).

Multimodal Representation Potential

Phylogenetic frameworks accommodate heterogeneous data types through layered visual encodings:

1. Branch thickness: Represents rhythmic complexity indices (e.g., syncopation density, hemiola frequency).

2. Node color gradients: Encode modal characteristics using Hue-Saturation-Value (HSV) mappings (e.g., Mixolydian mode = 240° hue, Dorian = 120°).

3. Leaf annotations: Attach cultural metadata (geographic distribution, associated dances) via interactive tooltips.

This multimodal approach bridges ethnomusicological analysis with computational rigor, enabling simultaneous examination of melodic evolution, cultural diffusion patterns, and performative contexts (Liebman et al., 2012).

Chapter 3

Prototyping and Co-Design Process

Taking into account all of the previously assessed information, this chapter details the process of designing and refining a visualization tool developed for this project. It begins with an introduction **Section** 3.1, and then the creation of initial mockups **Section** 3.2, including the thought behind design choices and the role of collaboration with this project's supervisor. The chapter then outlines the structure and execution of a co-design session with researchers/ experts **Section** 3.3, including the materials and script used. After going over the co-design session itself, it presents a synthesis of the feedback gathered during the session and discusses how these insights informed the subsequent development of the tool **Section** 3.4. Finally, this chapter goes over the experiment's limitations **Section** 3.5, and how they were managed throughout the process.

3.1 Introduction

The development of digital tools for musicological research, as all interfaces, relies on iterative design processes and close collaboration between researchers and domain experts. In the context of this project, the creation of a visualization tool required not only technical innovation but also a deep understanding of the needs and perspectives of its intended users. By combining insights from existing research with hands-on prototyping and participatory co-design methods, it was possible to create interfaces that are both functional and meaningful for research. This chapter explores the journey from initial mockup development to the integration of expert feedback through co-design, highlighting the value of collaborative and user-centered approaches in shaping digital research tools.

3.2 Initial Mockup Development

To envision the final product, we began by creating a series of mockups that reflected both the conducted research and the supervisor overview. Working with preliminary data from the project, rather than only mock data, allowed for the anticipation of practical challenges and ensured that early design decisions were grounded in real-world examples. This collaborative process helped align the prototypes with the project's objectives and the needs of potential users.

3.2.1 Early Prototypes

The first medium that was used to express our ideas was through paper drawings. Several sketches were drawn, and iterated upon weekly discussion with both of this project's supervisors (Figures 3.1 and 3.2).

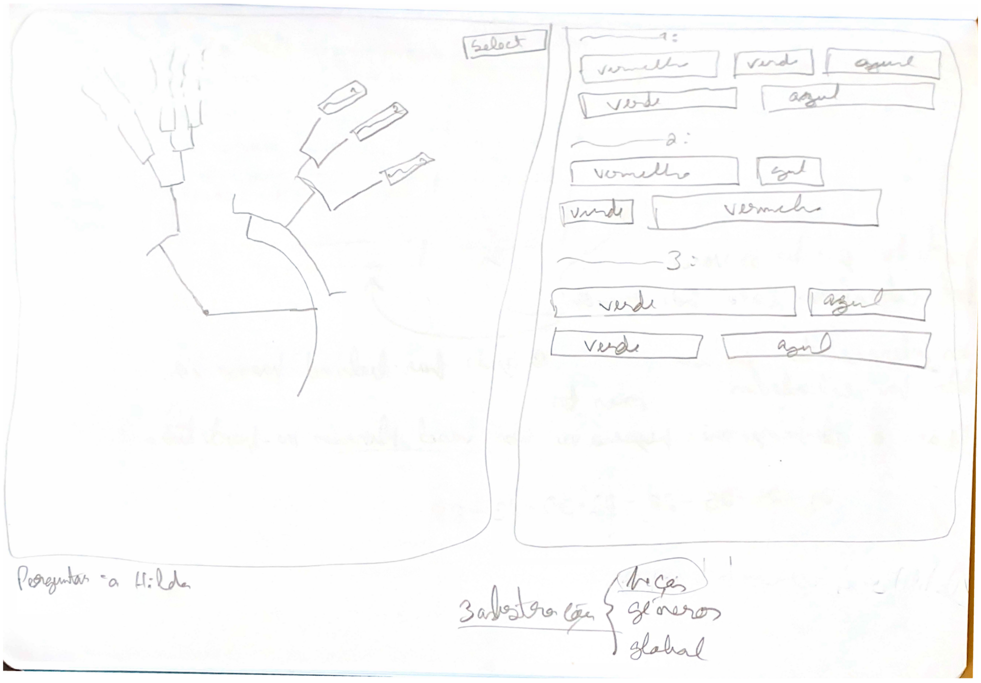

Figure 3.1: An initial sketch exploring potential applications of phylogenetic trees, including their use cases, and some notes with feedback from the supervisor

Once we had a clearer understanding of the project's development, we transitioned to using *Figma*, a widely-used and free prototyping software with which we had prior experience from other projects. This tool enabled the creation of a more precise and refined version of the prototype (Figure 3.3).

3.2.2 Prototypes

The initial prototypes were developed collaboratively, leveraging phylogenetic tree data files provided by this project's second supervisor (Figures 3.4). This research project applies bioinformatics methodologies, including sequence alignment and clustering techniques, to analyze similarities between folk music scores (in this project's scope, Irish and Galician traditions) across multiple levels:

Feature based: Chromatic (pitch content) and rhythmic patterns.

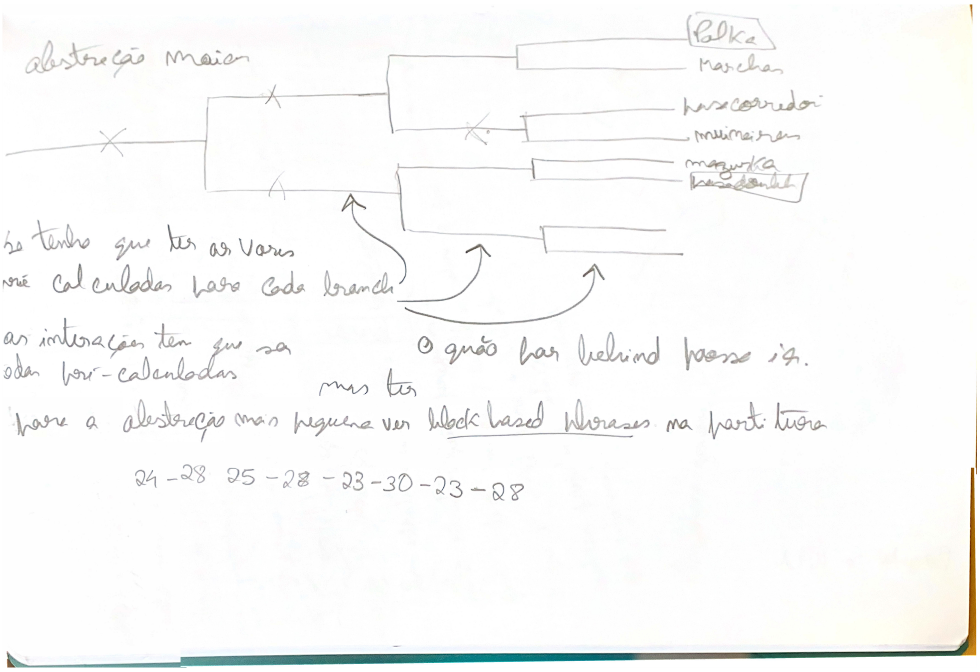

Figure 3.2: Another sketch exploring potential applications of phylogenetic trees, with some notes with supervisor feedback.

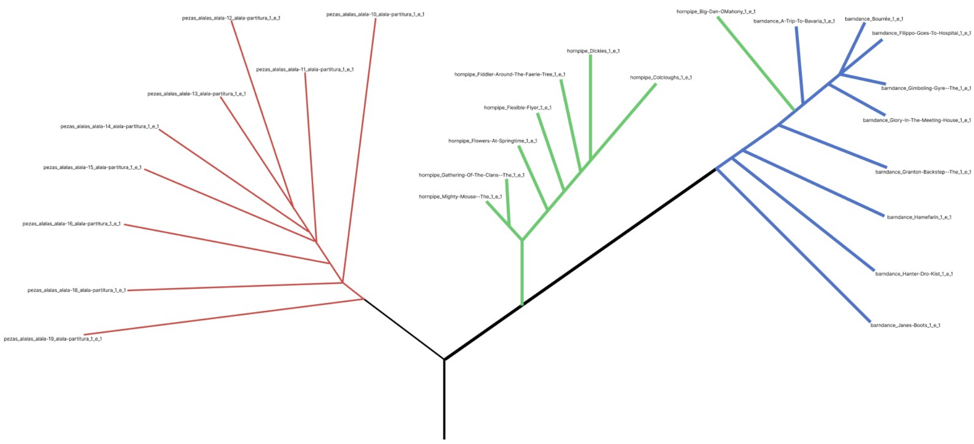

Figure 3.3: One of the first tests that was conducted on figma with phylogenetic trees. This tree does not contain real data, but it was foundational for me to understand the phylogenetic tree's structure.

Level Based: Note-to-note comparisons, shared melodic segments, formal structures (sequence of segments), and hybrid approaches.

Phylogenetic trees, were employed to visualize relationships between scores, genres, and traditions. The project supplied data in Nexus format, a standardized file structure for phylogenetic analysis.

Four iterative stages of this process are illustrated in Figures 3.4, 3.5, 3.6 and 3.7. Over several

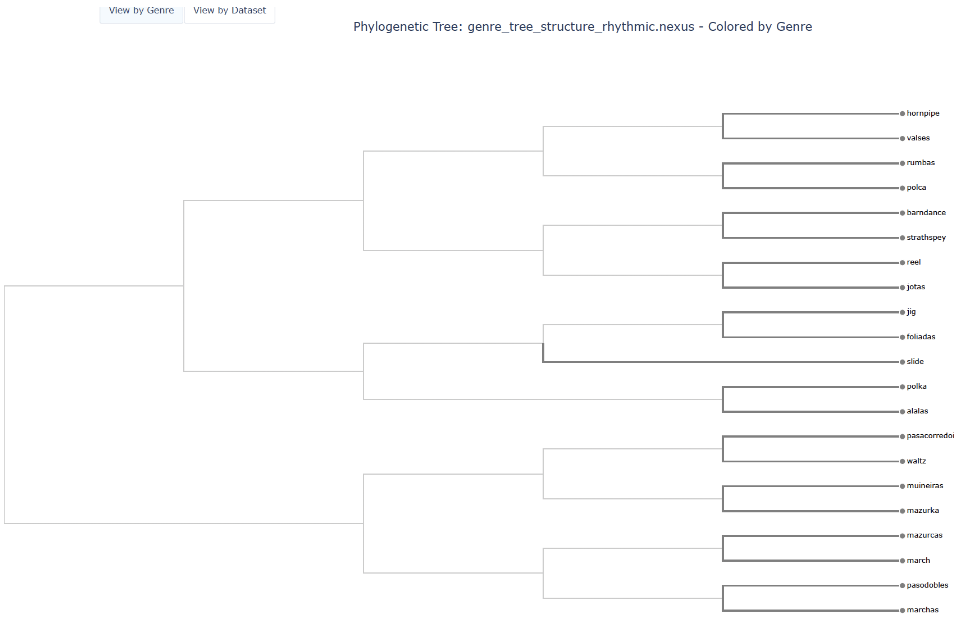

Figure 3.4: One of the initial tree data files was provided by the second supervisor. This file maps the relationships among all genres within both the Galician and Irish musical traditions, organizing them according to rhythmic features.

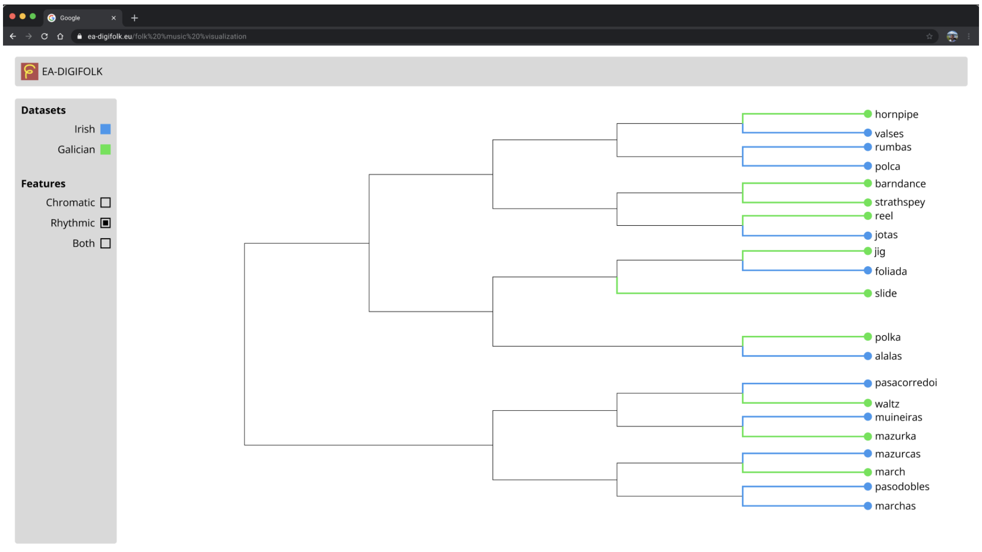

Figure 3.5: A direct translation of the file seen on (Figure 3.4).

weeks, the designs were refined through feedback loops with this project's supervisors, ensuring the visualizations balanced analytical rigor with usability for the target audience.

Upon the completion of several of these panels, the mockup was ready to be evaluated and scrutinized by some of the researchers who are directly working within this project's context.

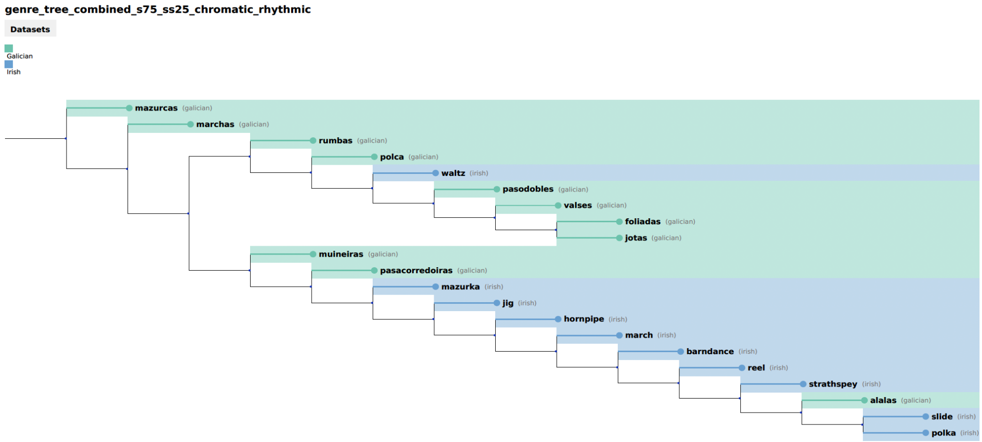

Figure 3.6: One more tree data files was provided by the second supervisor. This file maps the relationships among all genres within both the Galician and Irish musical traditions, organizing them according to chromatic and rhythmic features.

Figure 3.7: A direct translation of the file seen on (Figure 3.6).

3.3 Co-Design Session

In order to ensure that the visualization tool developed in this project meets the practical needs of musicologists and domain experts, a co-design session was conducted as a central component of the research methodology. This session was guided by several key questions:

1. What specific aspects of the visualization tool, such as clarity, usability, and research relevance, should be assessed to validate its effectiveness?
2. How do different visualization approaches, namely clustering and phylogenetic trees, compare in supporting scholarly analysis, both within a single genre and across multiple genres?

3. What challenges might participants encounter when interacting with computational representations of music, and how does the presence or absence of certain metadata influence the tool's research value?

4. What improvements or features do experts suggest to enhance the tool's analytical utility and scalability?

To address these questions, an online co-design session was organized, building on the initial prototypes developed for the project. The session utilized a series of interactive mockups and guided tasks (see A and B), and began with an introduction to the project's context and the specific goals of the co-design activity. The mockups, informed by research and supervisor input, served as the foundation for participant exploration and feedback.

Four scholars, all directly involved with the EA-DIGIFOLK project and possessing in-depth knowledge of the relevant datasets, participated in the session. Their expertise enabled them to provide particularly meaningful feedback. The session was intentionally interactive and reflective, encouraging participants to share their insights as they navigated the interface.

Recognizing that the logic behind phylogenetic representations might not be immediately intuitive, the session began with a brief explanatory introduction. The evaluation followed a guided format, with the moderator assigning specific, structured tasks to each participant.

The experimental design featured four distinct interaction flows, organized into two logical groupings: Genre-Specific Visualization and Multi-Genre Visualization. This structure allowed for a systematic comparison of visualization approaches and provided a robust foundation for analyzing participant feedback in relation to the guiding research question.

First Task: Genre-Specific Visualization

This pair of flows focused on evaluating visualization approaches for music within a single genre ("strathspey" was selected as the test case). Both flows followed identical initial steps, diverging only at the genre interaction stage:

Flow 1 presented the data in a clustering view (Figure 3.8), while Flow 2 displayed the same information using a phylogenetic tree structure (Figure 3.9).

Second Task: Multi-Genre Visualization

The subsequent flows examined how both visualization techniques performed with higher information density. This evaluation used a phylogenetic node containing six distinct genres: "strathspey", "polka", "slide", "alalas", "reel", and "barndance".

Flow 3, Clustering view of all six genres (Figure 3.10), while Flow 4, Phylogenetic tree representation of the same genre set (Figure 3.11).

Figure 3.8: The clustering view on the first flow. It showcases how it would like to see all the music on the strathspey genre with that philosophy.

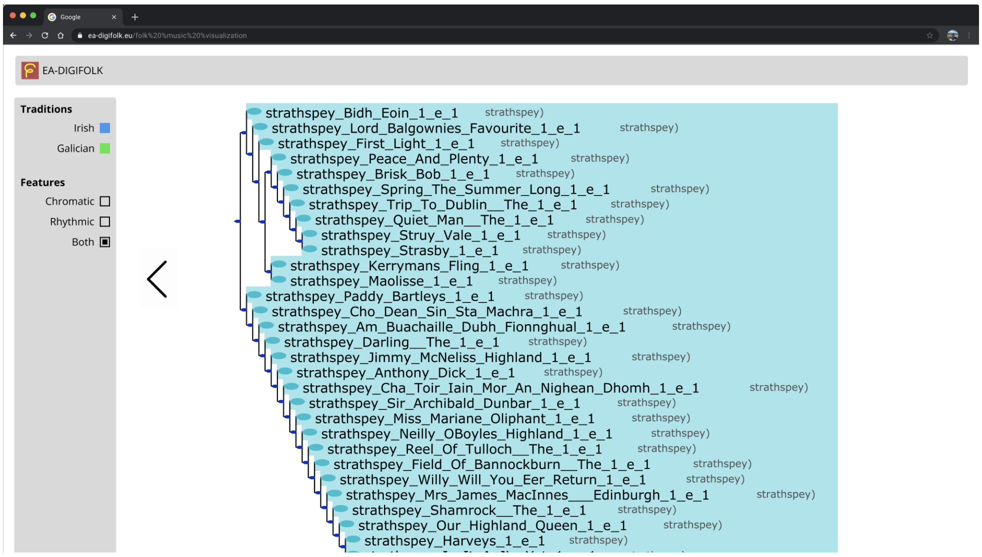

Figure 3.9: The phylogenetic tree view on the second flow. It showcases how it would like to see all the music in the strathspey genre with that philosophy.

3.4 Feedback

The controlled testing sessions with four domain experts (designated Participants 1–4) yielded substantive feedback across three critical dimensions: interface usability, musical representation adequacy, and system scalability. These insights fundamentally informed subsequent design iterations.

Participant 2 articulated a preference towards cluster-based visualizations over phylogenetic structures, describing the former as more immediately intuitive for melodic analysis. This perspective was tempered by Participant 1's observation that phylogenetic navigation became more

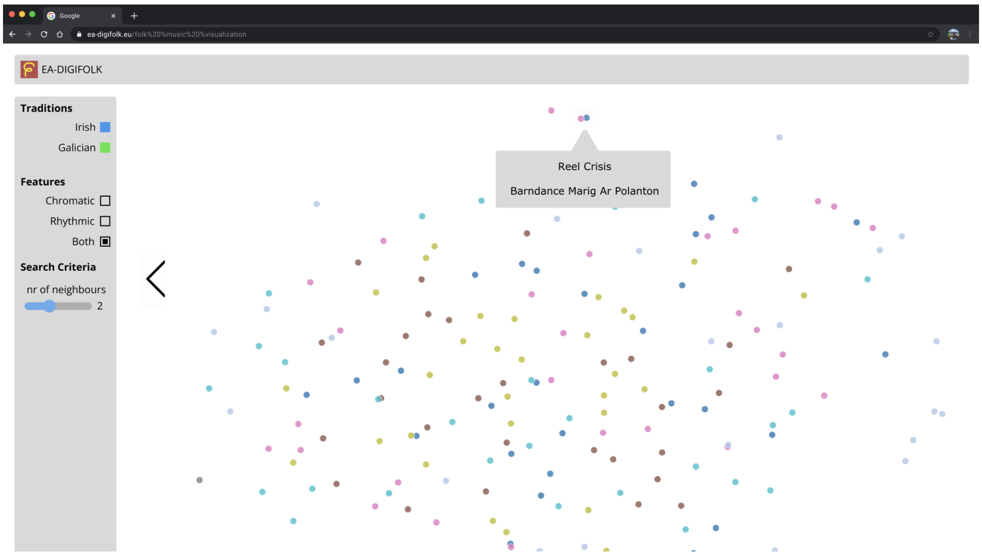

Figure 3.10: The clustering view on the third flow. It showcases how it would like to see all the music on the strathspey, polka, slide, alalas, reel and barndance genres with that philosophy.

Figure 3.11: The phylogenetic tree view on the fourth and last flow. It showcases how it would like to see all the music on the strathspey, polka, slide, alalas, reel and barndance genres with that philosophy.

intuitive with further exposure, suggesting the learning curve could be mitigated through improved initial instructions. Both musicologists nevertheless identified significant gaps in how musical knowledge was represented, with Participant 2 particularly emphasizing the absence of dance-related metadata, which isn't present in the datasets, but it was discussed during the session.

The musicologists' feedback highlighted a fundamental tension between computational abstraction and musicological practice. Participant 1's repeated references to phrase structures, notably the A-B-A form (a three-part musical form consisting of an opening section (A), a following section (B) and then a repetition of the first section (A)), and Participant 2's insistence on movement notation revealed that the current feature set, while computationally robust, failed to capture essential analytical categories used in ethnomusicology. Their collective suggestion to incorporate a similarity quantification metric, potentially through heatmap annotations, points toward several necessary enhancements to the visualization.

Participants 3 and 4 provided critical technical corrections regarding fundamental representation issues. Participant 3's analysis of UMAP's dimensional reduction artifacts demonstrated how spatial proximity in the visualization frequently misrepresented actual musical relationships, a problem exponentially compounded when applied to Ireland's dataset of more than twelve thousand scores. Participant 4 extended this critique by demonstrating how the interface's attempt to render cross-traditional comparisons created visual noise that undermined analytical utility.

The two domain experts (Participants 3 and 4) agreed that phylogenetic trees offered the most reliable framework for representing musical relationships, but identified two key improvements needed. First, Participant 3 stressed that the visual representation must clearly communicate how connection strengths between musical pieces are calculated, for instance, through distinct line weights, colors, or interactive labels explaining relationship metrics. Second, Participant 4 suggested these trees should serve primarily as organizational frameworks to guide users, who could then access more specialized analytical views (like heatmaps) for deeper investigation of selected musical clusters. Together, their feedback pointed toward a combined visualization strategy: using trees for initial orientation and navigation, while providing dynamic, user-controlled tools for detailed comparison of specific musical elements.

3.5 Conclusion

Synthesizing the feedback from the co-design session in light of the guiding research questions, three core design imperatives emerged that directly address the effectiveness, comparative value, challenges, and future enhancements of the visualization tool.

First, in assessing clarity, usability, and research relevance, participants highlighted the importance of progressive disclosure. Participant 4, in particular, advocated for an interface that clearly showcases inter-traditional relationships through filtered, focused views before allowing users to expand the scope of their analysis. This approach directly combats the information overload observed during testing, ensuring that users can engage with the most meaningful comparisons from the outset.

Second, when comparing clustering and phylogenetic tree visualizations across single, and multi-genre contexts, feedback underscored the need for user-driven control over visible elements. Participant 3's suggestion for user-chosen element rendering offers a practical solution to scalability, by allowing users to dynamically adjust what is displayed based on their interests and device capabilities, the tool can remain both powerful and accessible, regardless of dataset size or complexity.

Third, the session revealed key challenges in user interaction, particularly the gap between computational representations and musicological practice, as well as the impact of missing metadata. Participants unanimously called for embedded guidance features, ranging from references to established visualization models, to interactive heatmaps, to help users interpret complex structures and make informed analytical decisions.

Collectively, these insights point to a necessary shift from comprehensive, all-encompassing representations to focused, contextually meaningful views. While the phylogenetic tree stands out as the optimal structural framework for comparative musical analysis, it must be enhanced with contextual metadata, real-time analytical tools, and adaptive visual complexity. Such refinements will not only support rigorous scholarly inquiry but also maintain accessibility for musicologists, effectively addressing the core questions that shaped the co-design session and guiding the next phase of tool development

3.5.1 Experiment's limitations

While this study provided valuable insights into participants' understanding and interaction with the proposed visualization system, several limitations must be acknowledged. Due to time constraints and the researcher's intermediate proficiency with Figma, certain aspects of the prototype could not be fully developed.

A notable example is the clustering view of the five genres (strathspey, polka, slide, alalas, reel, and barndance) presented in the third experimental flow. While the clustering logic itself accurately reflected the underlying data, its visual implementation in the prototype was suboptimal, potentially misleading participants about its efficacy as a representation method. During testing, the researcher clarified that the final implementation would differ significantly, but the absence of a more refined visual demonstration remains a limitation, one that may have influenced participants' perceptions of this approach.

These constraints highlight the need for future iterations to: Improve prototype fidelity, particularly for complex visualizations like multi-genre clustering, allocate more development time for high-priority interactive elements, and conduct pilot testing to identify and address usability issues before formal evaluation.

Chapter 4

Prototype Implementation

This chapter presents the technical journey behind the development of the folk music visualization tool, tracing the path from conceptual foundations to a fully functional web-based system. It opens with an introduction **Section** 4.1 that situates the implementation work within the broader research goals, establishing the conceptual framework for visualizing structural relationships in folk music. Following this, the architecture overview **Section** 4.2 details the key design decisions and system structure, explaining how these choices enable responsive, scalable, and maintainable visualizations. **Section** 4.3, describes the musical data sources, preprocessing workflows, and annotation strategies that transform raw musical material into structured, analyzable data for visualization. The implementation **Section** 4.4 provides an in-depth account of the development process, covering the integration of visualization components, user interface design, and performance optimization techniques that together form a cohesive analytical tool. Finally, section **Section** 4.5 presents the co-design session followup, analyzing expert feedback on the implemented system and explaining how this feedback informed refinements to better serve the need of ethnomusicology researchers. Throughout the chapter, we highlight the iterative development process that bridges theoretical musicology concepts with practical visualization techniques.

4.1 Introduction

The transition from the theory to functional implementation was a critical phase in this research. With established musicological frameworks and validated user requirements from expert evaluations, the project faced its next fundamental challenge: selecting and deploying a technical architecture capable of realizing the envisioned visualization tool. The mandate for a web-based solution necessitated careful consideration of weighing factors ranging from computational efficiency to long-term maintainability within the broader EA-DIGIFOLK research ecosystem.

After consultation with both this project's supervisor and EA-DIGIFOLK technical partners, **Vue.js** was pointed as the optimal framework. As a progressive JavaScript solution, it combined the responsiveness required for dynamic musical data visualization with the accessibility needed for interdisciplinary collaboration. Its existing adoption across related project tools further ensured coherence between the whole project. This decision proved particularly smart given JavaScript's diverse ecosystem, although there is no singular universal solution is ideal for this type of project.

What follows documents this technical journey: how phylogenetic representations of musical relationships assumed functional form, how interface elements addressed the usability concerns raised during testing, and ultimately, how a working tool emerged to serve both analytical and exploratory research purposes.

4.2 Technical Workflow

To prepare the phylogenetic trees for web visualization, it was developed a conversion process that transforms NEXUS files (the standard format for phylogenetic analyses) into JSON (JavaScript Object Notation) format, which is readily compatible with modern web visualization libraries.

The conversion process begins by using the "BioPython" library to read and parse each NEXUS file, extracting the hierarchical structure that represents evolutionary relationships between musical pieces. During transformation, the system traverses each phylogenetic tree, building a nested JSON structure that mirrors the original hierarchy. For leaf nodes (which represent individual musical pieces), the process extracts and cleans music's names, removing extra characters and standardizing formats. The system also identifies and extracts genre information where available from the music's names, enabling more sophisticated categorization in later visualization stages.

Before generating the final output, each converted tree undergoes validation to ensure proper structure. This validation checks for required fields, proper nesting of child nodes, and other structural requirements that ensure compatibility with visualization libraries. When issues are detected, warning messages highlight potential problems that might affect visualization quality.

The final JSON files maintain the structural integrity of the original phylogenetic data while transforming it into a format that web-based visualization tools can easily process. This conversion forms the foundation for subsequent processing steps, where the data is further refined and segmented according to musical traditions and genres. These additional segmentation enable focused analysis of relationships within specific musical contexts, providing deeper insight into how musical forms have evolved within their cultural traditions.

4.3 Architecture Overview

The application employs a **component-based architecture**, where specialized modules manage separate functional aspects of the system. The core components consist of a visualization area, which acts as the main interface for presenting phylogenetic trees, and a sidebar, dedicated to filtering data and adjusting visualization settings. Within the visualization area, a score viewer renders "MusicXML" files, and a breadcrumb navigation system supports hierarchical exploration of the data.

The component architecture was implemented using Vue.js's composition API (Application Programming Interface), which provides a more flexible approach to organizing code. This choice facilitated better separation of concerns, with each component maintaining its own internal state while communicating with other components through clearly defined interfaces. The reactive state management in Vue.js was particularly valuable for maintaining consistency across the different views as users navigate through the hierarchical data.

The system maintains its state through reactive variables that track the current view, selected tradition, genre, and music score. This state management approach enables seamless transitions between different levels of the data hierarchy. When state changes occur, the system efficiently updates only the components affected by those changes, preserving performance even when handling complex visualizations.

As illustrated in (Figure 4.1), the system architecture operates across four distinct layers, each serving specific functional responsibilities. The Component Layer houses the primary user interface elements, including the visualization area component that renders phylogenetic trees using "D3.js", the sidebar component for parameter control, and breadcrumb navigation for hierarchical exploration. Below this, the Visualization Layer integrates specialized libraries: "D3.js" for interactive tree visualization and OpenSheetMusic Display (OSMD) for score rendering.

The Data Processing Segment constitutes the system's computational core, overseeing four critical services. Tree processing functions ("enhanceTree", "filterTree", "renderTree") manage phylogenetic data manipulation and rendering pipeline, while score processing modules ("mapScoreToFilePath", "loadScoreContent", "RenderWithOSMD") handle "MusicXML" file operations. Navigation services facilitate seamless transitions between data hierarchy levels through functions like "navigateToBreadcrumb" and "LoadGenre Phylotree", supported by file path resolution utilities that dynamically locate and validate data resources across the corpus structure.

The Data Layer provides persistent storage through organized JSON file containing traditions and genre trees alongside Irish and Galician "MusicXML" collections. This layered architecture ensures scalable data handling while maintaining responsive user interactions across the phylogenetic exploration interface.

4.4 Dataset and Similarity Annotations

The visualization's filtering system evolved considerably during development. Initial mockups used temporary terms that reflected some uncertainty about feature definitions, but after reviewing EA-DIGIFOLK research and testing different approaches, we settled on three core musical features: "Chromatic", "Rhythmic", and "Chromatic & Rhythmic" form. While we use the technical term "Features" in this project, musicologists might better recognize these as "Musical Features", a distinction that helps connect our technical system with traditional music analysis methods.

The features themselves follow an established music analysis framework, adapted for this specific project. The Chromatic Interval measures the distance between notes in semitones (where B to C equals 1 semitone), with negative numbers showing descending movement (Romero-Velo et al., 2025). This allows comparison of melodies regardless of their key. The Rhythmic Ratio compares note durations (like a quarter note followed by a dotted quarter note, giving a ratio

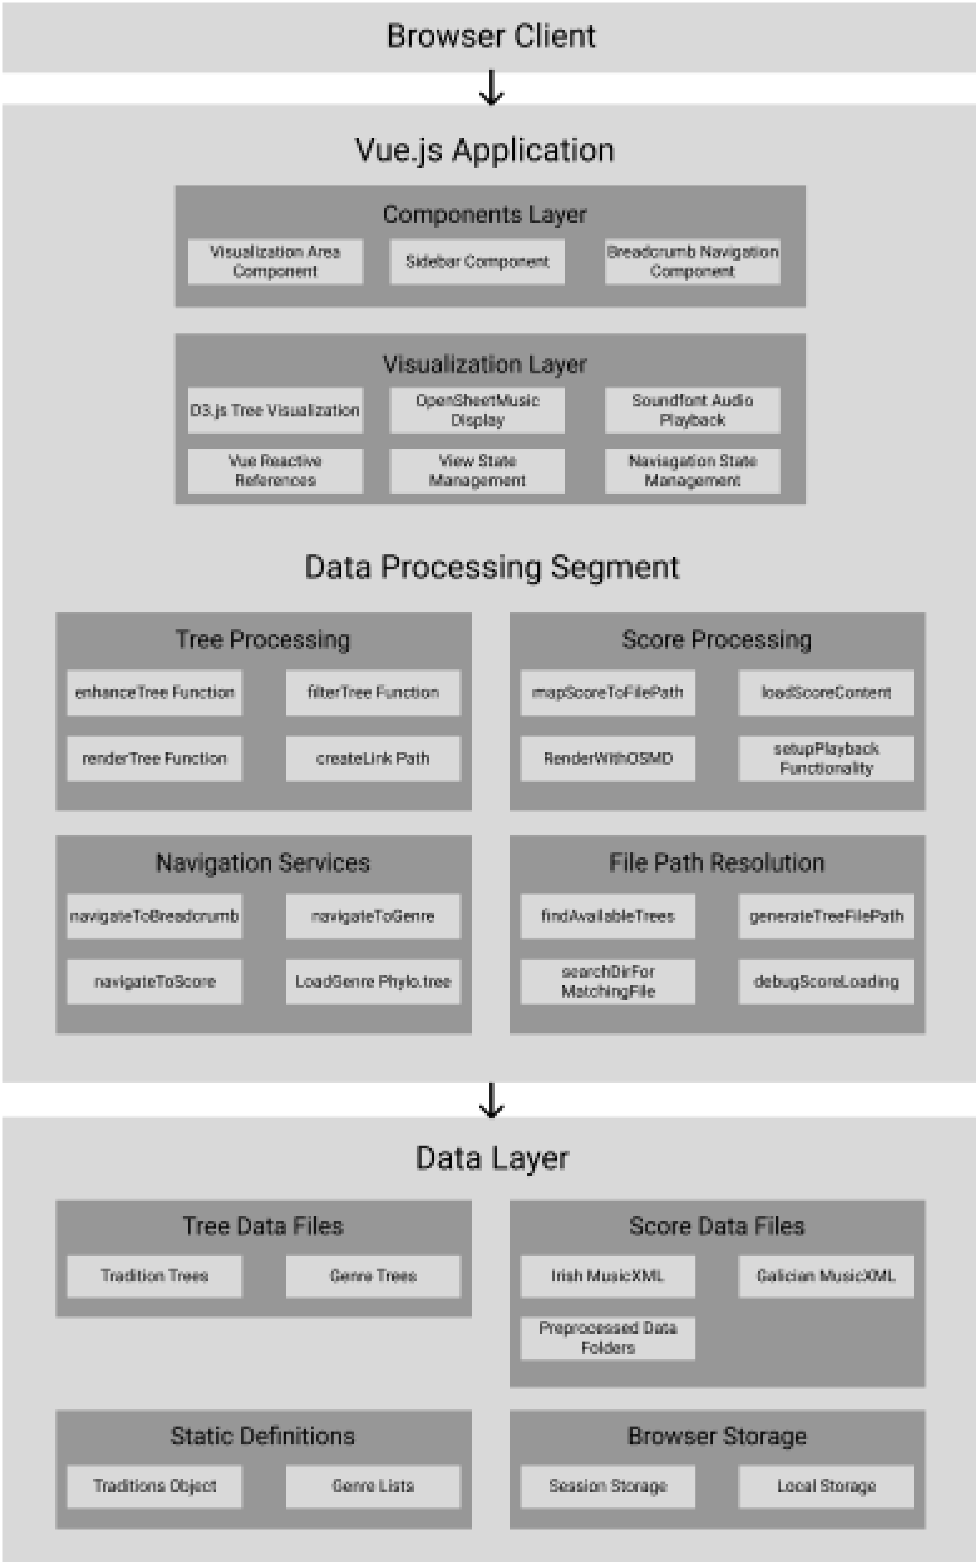

Figure 4.1: This is the proposed computational architeture websystem for this project.

of 1.5), and carefully handles rests as meaningful musical events rather than silences to ignore (Romero-Velo et al., 2025). The combined Chromatic-Rhythmic feature represents these two aspects together in a single representation that shows both pitch and rhythm patterns simultaneously (Zhu et al., 2022; Romero-Velo et al., 2025).

We implemented four levels of musical comparison in the visualization. The simplest, Note-Level comparison (called Global Similarity in technical terms), that looks at entire melodies side-by-side. The Shared Phrases method identifies recurring musical patterns within pieces, while the Form approach examines how these patterns are arranged over time. These combine in our final and most sophisticated analysis (called Combined), which blends both pattern recognition and structural analysis with a 75%-25% weighting favoring form (Heyer et al., 1999; Needleman and Wunsch, 1970; Romero-Velo et al., 2025).

This carefully designed system serves two important purposes. First, it focuses on the most musically meaningful features for comparing folk dances, setting aside less relevant technical details. Second, it maintains crucial elements like rhythmic accuracy that are essential when studying dance music. The result is a visualization tool that is both computationally sound and genuinely useful for music researchers.

4.5 Implementation

After establishing the technical architecture and data processing strategies, the next stage of development focused on constructing the main structure of the visualization tool. As previously mentioned, to ensure a coherent and user-friendly interface aligned with the initial prototypes, the system was organized into two primary segments the main visualization window and the filters section. The filters are divided into two groups, “Features” (Chromatic, Rhythmic, and Chromatic & Rhythmic) and “Level” (Note, Shared Phrases, Form, and Combined)—with users able to select one option from each group to tailor their exploration (the visualization is available to consult at pelinhx.github.io/eadigitest).

The main visualization area is designed to support navigation across several analytical abstractions that mirror the hierarchical nature of the dataset. Users begin at the Tradition phylogenetic tree, which offers a broad overview of musical lineages. From there, they can delve into the Genres tree to examine stylistic categories within each tradition, then further into the music tree to explore relationships between individual pieces, and finally access either the score, or the comparison view for detailed side-by-side analysis of selected works. This layered approach not only reflects the structure of the underlying data but also supports both broad exploration and focused analytical tasks, ensuring that the tool remains accessible and effective for diverse research needs within ethnomusicology.

4.5.1 Navigation and User Interface

The user interface was designed to support exploratory research through intuitive navigation. A breadcrumb system tracks the user’s position within the hierarchical data structure, allowing for easy movement between different levels of abstraction. Traditions view (the abstraction’s top level) (Figure 4.2), tradition's genre view (Figure 4.3), music relation view (displaying phylogenetic relationships between pieces in a genre) (Figure 4.4) and individual score view (showing the musical notation for a selected piece).

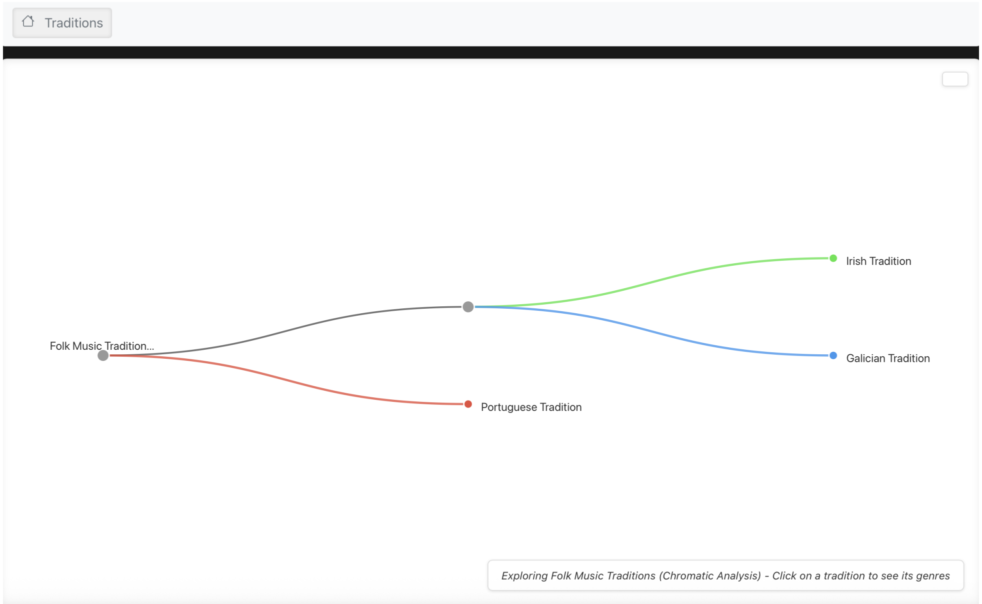

Figure 4.2: The first level abstraction when you open up the visualization tool. It features a very simple, hardcoded, phylogenetic tree with a very simple layout, just to fit as an example of what could be an even higher abstraction layer with many more traditions.

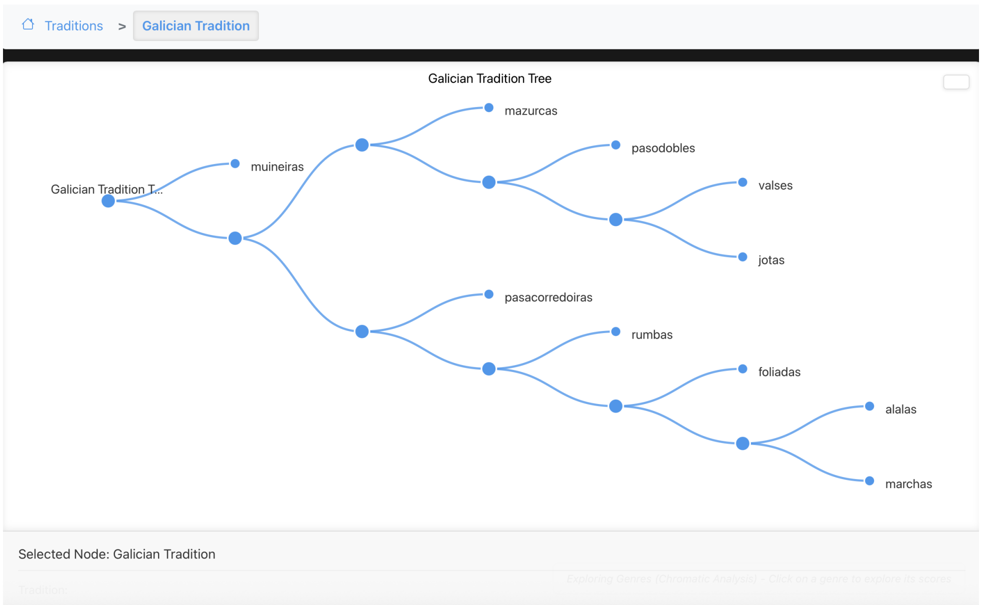

Figure 4.3: This is the phylogenetic tree corresponding to all the music genres in the Galician Tradition.

The breadcrumb implementation maintains the complete navigation path, with each segment representing a level in the data hierarchy. This not only provides users with context about their

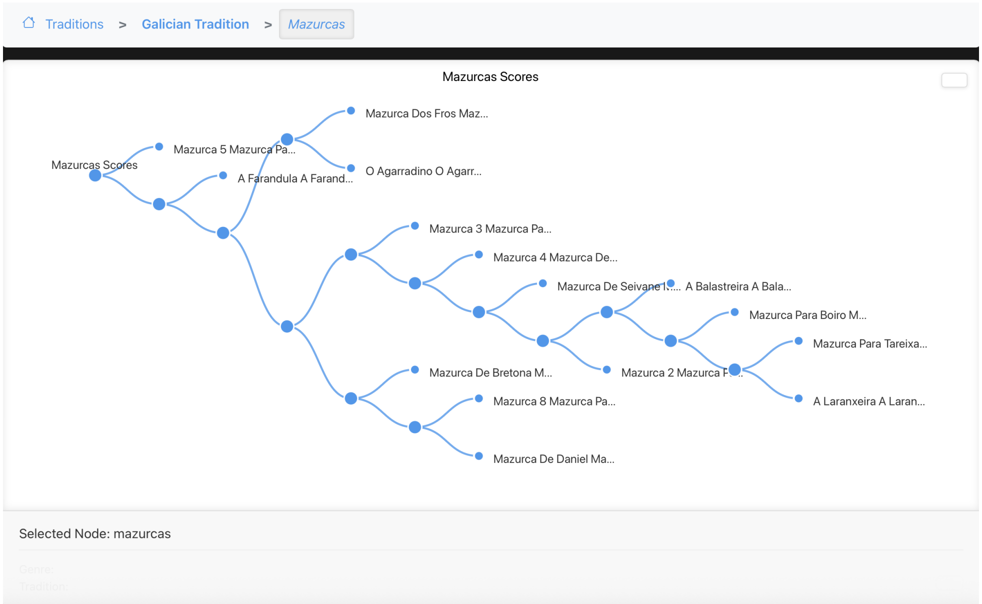

Figure 4.4: This is the phylogenetic tree corresponding to all the music within the "Mazurcas" genre.

current location but also serves as an interactive navigation tool. Each breadcrumb element is clickable, allowing users to move directly to any previously visited level without having to step backward sequentially (Figure 4.5). This bidirectional navigation capability significantly enhances the exploratory potential of the application.

Figure 4.5: Example of the breadcrumb implementation, in this case it features the path through "Irish Tradition", on the "Barndance" genre, and onto the song "Amaryllis Schottische The 1 E 1"

The interface adapts to the current view, providing context-appropriate controls and information. This design facilitates both broad exploration of musical relationships and detailed examination of specific pieces. Visual consistency is maintained across different views through a cohesive color scheme and interaction pattern, helping users build a mental model of the data structure as they navigate through it.

4.5.2 Data Organization

The data structure follows a hierarchical organization that mirrors the categorizing nature of folk music traditions. At the highest level, music is categorized by tradition (Irish, Galician), followed by genres specific to each tradition (such as "jigs", "reels", and "hornpipes" for Irish music, and "muiñeiras", "alalas", and "jotas" for Galician music). Each genre contains individual musical pieces represented as leaf nodes in the phylogenetic trees.

The data's folder structure follows this hierarchy, with separate directories for each tradition and subfolders for genres. The preprocessing of musical data generated specialized JSON files that capture the phylogenetic relationships between pieces based on different analytical approaches. These files are organized according to both tradition and genre, with naming conventions that reflect the specific features and analysis levels used to generate them.

The phylogenetic tree data itself uses a nested JSON structure that captures parent-child relationships, with additional metadata about musical attributes attached to each node. This organization allows for efficient traversal and rendering of the hierarchical data while maintaining the rich contextual information needed for meaningful musical analysis.

4.5.3 Visualization Implementation

The core visualization functionality in the EA-DIGIFOLK project is implemented using D3.js to create interactive phylogenetic trees representing relationships between folk music pieces. The implementation employs several key optimization techniques to handle the rendering and interaction with complex hierarchical data.

The system implements adaptive rendering based on tree density through the "getContentAwareLayoutConfig" function. This function dynamically adjusts visual parameters like node size, text visibility, and spacing based on the number of nodes in the tree. For denser trees with many leaf nodes, the system automatically reduces node radius, adjusts text visibility thresholds, and modifies branch thickness to maintain readability while preserving performance.

The rendering pipeline adapts its behavior based on the current view context (tradition, genre, score), providing appropriate visual emphasis to different elements depending on the navigation level. This context-aware approach ensures that the most relevant information is highlighted at each level of the hierarchy.

A form of viewing frustum is implemented within the zoom handler, where text labels are conditionally rendered based on zoom level. This technique prevents performance degradation when visualizing dense trees with numerous leaf nodes (more than 30 leaf nodes, in this case) by only displaying text for leaf nodes when the user has zoomed in sufficiently. This selective rendering approach is crucial for maintaining responsive performance when dealing with complex phylogenetic relationships.

The system distinguishes between internal nodes (representing genres or traditions) and leaf nodes (representing individual scores) with different visual treatments. Internal nodes are rendered smaller than leaf nodes at the extremities, creating a clear visual hierarchy that enhances navigation while conserving rendering resources for the most relevant content.

The visualization implements sophisticated path resolution through the "findAvailableTrees" function, which detects the current environment, implements cascading fallback paths, and caches successful paths in session storage to reduce redundant network requests. This approach ensures robust loading even when file locations are inconsistent or when deployed to different environments.

Several performance optimization techniques are employed throughout the codebase, including conditional ”DOM” (Document Object Model) updates wrapped in ”nextTick()” to prevent ”UI” (User Interface) blocking, sequential loading for score comparison to prevent browser crashes, smart error recovery with fallbacks when expected data isn’t available, and delayed text collision detection to prevent label overlap in dense areas.

These techniques work together to provide responsive performance even when visualizing complex phylogenetic relationships with hundreds of nodes, demonstrating a sophisticated approach to visualization that balances visual fidelity with performance considerations for dense hierarchical structures.

4.5.4 Score Rendering System

A significant challenge in the implementation was the integration of musical score rendering. This was accomplished by incorporating the “OpenSheetMusicDisplay” library, which interprets “MusicXML” files and renders them as standard music notation. The score rendering system includes, dynamic loading of “MusicXML” files based on selected musical pieces, validation of “MusicXML” content to ensure proper rendering, fallback mechanisms for handling malformed or missing files and download functionality for researchers to access the original files.

The implementation of “OpenSheetMusicDisplay” required careful consideration of asynchronous loading patterns. Rather than loading the library for all views, it employs a dynamic import strategy that only loads the rendering engine when a user actually navigates to a score view. This approach optimizes initial load times while ensuring that the full score rendering capability is available when needed.

The score renderer provides a structured approach to displaying musical content, with a clear header showing title and composer information extracted from the “MusicXML” metadata. The rendering process itself is wrapped in comprehensive error handling that defaults to alternative representations when issues are encountered, such as showing the raw XML with syntax highlighting if the rendering engine cannot properly interpret a particular score.

The implementation accommodates differences in file naming conventions between traditions, employing a flexible path resolution system that tries multiple potential file locations when the primary path fails. This path resolution system includes fallback strategies that attempt variations of file names based on known patterns specific to each tradition, significantly improving the robustness of the score retrieval process.

4.5.5 Score Comparison

A key feature of this visualization tool is the score comparison functionality, which enables researchers to examine multiple musical pieces simultaneously to identify structural patterns across related works. This feature directly addresses a core need identified during co-design sessions, the ability to visually analyze similarities between folk music compositions.

The implementation follows a sequential workflow beginning when users activate comparison mode from the genre view. Users can select up to three scores from the phylogenetic tree, with this limit established based on interface constraints and cognitive load considerations. The system provides visual feedback through node highlighting to indicate selected scores before transitioning to a dedicated comparison view where each selection is rendered side-by-side in a responsive grid layout.

The score rendering leverages ”OpenSheetMusicDisplay” to transform “MusicXML” data into interactive notation. To prevent browser performance issues, the system implements sequential rather than parallel loading of scores. Each comparison panel includes individual zoom controls, allowing researchers to focus on specific sections while maintaining overall context. The tool also provides download functionality for each score, enabling researchers to obtain the original “MusicXML” files for further analysis in specialized software.

A breadcrumb navigation system maintains context throughout the comparison process, allowing users to easily return to the genre view or explore other scores. This implementation demonstrates how specialized visualization tools can enhance traditional musicological analysis by facilitating direct comparison of related works within a larger corpus.

Chapter 5

Evaluation

This chapter delivers a critical evaluation of the folk music visualization tool, considering both its strengths and areas for further development. It opens with a follow-up on the co-design session **Section** 5.1, where expert feedback is gathered through a structured questionnaire to assess how well the implemented system aligns with initial expectations and addresses previously identified usability issues. **Section** 5.2, discusses the technical and design obstacles encountered during development, outlining the solutions and adaptive strategies that were implemented to ensure system robustness and usability. Finally, **Section** 5.3, presents a synthesis of expert evaluations, combining standardized usability metrics and qualitative insights to highlight validated successes and identify priorities for future refinement.

5.1 Co-design Session Followup

After all this work, the next step was to get expert opinions on the quality of the tool. Following the initial co-design session, conducting a follow-up evaluation with the same experts was crucial for for two main reasons reasons. First, it allowed validation of whether the implementation matched their expectations from the mockups, and it provided an opportunity to assess if the previously identified usability issues had been adequately addressed.

Rather than organizing another time-intensive in-person session, I opted for a more streamlined approach using a Google Forms questionnaire. This method offered several advantages: it accommodated researchers' busy schedules, allowed flexible participation, and reached a wider audience of relevant experts beyond the original co-design participants.

The evaluation was structured in three distinct sections to gather comprehensive feedback.

5.1.1 Introduction and Context

The first section provided necessary background information about the visualization tool and its purpose. It explained the concept of phylogenetic trees in the context of folk music, describing how branches represent evolutionary relationships between pieces, nodes represent scores or related music groups, and distance indicates musical similarity. This introduction was essential for participants who hadn't been part of the initial co-design session, ensuring they understood the conceptual framework before interacting with the tool (See Annex X).

Participants were given a link to the application and asked to spend around 5 minutes exploring the tool, with specific guidance to view scores from a genre of their choice and try the comparison feature with up to three scores.

5.1.2 Standardized Usability Assessment

The second section utilized the System Usability Scale (SUS), a standardized quantitative tool that offers an objective measure of usability. By employing SUS, we could compare the tool's usability against established benchmarks and track potential improvements in future iterations.

Since usability is context-dependent and lacks universal measurement standards, flexible and widely applicable assessment tools are essential. SUS meets this need by providing a quick, reliable method for evaluating system usability, enabling meaningful comparisons even in the absence of absolute benchmarks (Brooke, 1995).

> "To calculate the SUS score, first sum the score contributions from each item. Each item's score contribution will range from 0 to 4. For items 1,3,5,7,and 9 the score contribution is the scale position minus 1. For items 2,4,6,8 and 10, the contribution is 5 minus the scale position. Multiply the sum of the scores by 2.5 to obtain the overall value of SU. SUS scores have a range of 0 to 100." (Brooke, 1995)

Given this structured approach, SUS plays a crucial role in this project's usability assessment, providing valuable feedback for evaluation.

5.1.3 Qualitative Feedback

The final section featured open-ended questions designed to elicit detailed insights about the tool's functionality and potential applications.

For returning participants from the initial co-design session, additional questions explored how well the implemented visualization matched the discussed mockups and whether they encountered any unexpected navigation difficulties.

All participants were asked about practical applications for the tool in their work or research, the clarity and functionality of the score visualization, and any additional suggestions to make the tool more valuable.

This adaptive questionnaire design allowed for detailed feedback collection while maintaining a reasonable completion time, making it more likely that a larger number of experts would participate. By combining standardized usability metrics with open-ended qualitative responses, the evaluation aimed to gather both comparable numerical data and rich descriptive insights about the tool's effectiveness for its intended audience.

This approach recognized that ongoing feedback from domain experts is essential when developing specialized visualization tools, especially for niche fields like folk music analysis where the intended users possess specific domain knowledge and research practices that must be accommodated.

5.2 Overcoming Challenges

Being this a project that involved extensive coding, it presented several challenges in implementing all intended features. Considering the time constraints, the development process required careful prioritization and effective problem-solving. Several significant technical and design challenges were encountered throughout the implementation process.

As mentioned, one of the most persistent challenges was the handling of "MusicXML" files across different folk music traditions. The Irish and Galician datasets followed different naming conventions and organizational structures, requiring the development of a more sophisticated file path resolution system. This was solved through the implementation of a multi-tiered fallback system that attempts various path patterns when the primary path fails, significantly improving the robustness of score retrieval.

The phylogenetic tree visualization posed its own set of challenges, particularly when dealing with large datasets. Dense trees with numerous leaf nodes often became cluttered and difficult to navigate. This was addressed through the implementation of dynamic layout configurations that adjust spacing, node size, and text visibility based on tree density. Additional features such as adaptive zooming and contextual text hiding further enhanced the user experience.

Navigation between different levels of the data hierarchy required considerable attention to state management. The breadcrumb system initially failed to update correctly during certain navigation paths. This was resolved by implementing a more rigorous state update sequence, to ensure proper component re-rendering after state changes.

File path resolution across different deployment environments (development and production) presented a particularly hard challenge. This was solved by implementing a path discovery algorithm that tries multiple base paths and retains successful paths in session storage for future use, significantly improving loading reliability across environments.

Throughout the implementation process, these challenges necessitated constant refinement of the codebase, with particular attention to error handling and performance optimization. The solutions developed not only addressed the immediate issues but also contributed to a more robust and flexible system architecture that can accommodate future enhancements and dataset expansions.

5.3 Co-design Session Followup Feedback

The evaluation of the visualization tool involved fifteen participants who provided feedback through an online survey. The survey was created using Google Forms and circulated among both the previous session participants and several musicology-focused Google groups to gather diverse perspectives.

As previously mentioned, the survey was structured in two complementary sections. The first section employed the System Usability Scale (SUS), and the second one solicited detailed qualitative comments, allowing participants to express observations about the tool's functionality, usability challenges, and potential research applications.

For the first segment, we got the results on the Figure 5.1. Applying the aforementioned calculation for each of the participants and getting the average score between all, we get to a result of **69.33**. Although this is a good score, it shows that there is still some room for improvement.

Carimbo de d	I think that I v	I found the sy	I thought the	I think that I v	I found the va	I thought ther	I would imagi	I found the sy	I felt very con	I needed to le
2025/06/05 8	5	1	4	1	3	1	4	1	4	1
2025/06/05 1	5	4	2	3	4	4	1	3	4	5
2025/06/06 1	1	2	5	1	1	4	1	5	1	1
2025/06/06 6	4	2	4	3	4	1	4	3	3	2
2025/06/07 1	4	1	4	2	3	2	4	1	3	1
2025/06/08 1	2	2	3	1	3	1	2	3	3	2
2025/06/08 3	2	2	4	3	4	1	3	3	3	2
2025/06/09 3	3	2	3	4	4	3	5	3	3	4
2025/06/09 6	3	2	4	1	3	2	4	2	3	2
2025/06/11 1	4	1	5	2	5	1	5	2	4	5
2025/06/11 8	4	2	5	1	5	2	5	1	5	1
2025/06/12 3	2	1	5	1	3	1	4	1	2	3
2025/06/12 7	3	1	5	1	5	1	3	1	3	1
2025/06/16 4	4	1	5	1	4	2	4	1	4	4
2025/06/17 3	3	3	4	2	3	3	4	2	4	2

Figure 5.1: These are the results for all the questions on the SUS among the fifteen participants. Each column represents each of the ten questions, while the rows represent each of the participants.

On the second section, we got qualitative results, which, adding to the score we achieved on the SUS, will help us understand what we will need to implement on future versions of this visualization.

Participants offered detailed comments regarding both the strengths and limitations of the implemented visualization. Several respondents highlighted that the current version aligns well with the initial mockups and co-design discussions. Notably, users appreciated the tool's potential for supporting research tasks such as analyzing large collections of scores, identifying similarities between musical pieces, and facilitating comparative analysis across different traditions and genres.

However, a number of usability challenges and feature requests were identified, on the layout and design, some users suggested minor web design adjustments to optimize space utilization and improve the overall layout. There were also requests for a more "clean" and intuitive design, with particular attention to reducing visual clutter when displaying multiple songs or styles. For score visualization, it was generally described as clear and functional, especially for exploring structural relationships between pieces. Suggestions included the introduction of a search bar for easier navigation and the option to compare scores from different genres or traditions. On the functionality and features standpoint, users expressed interest in additional features such as the ability to upload and analyze custom datasets, support for audio playback, and the inclusion of contextual information (e.g., annotations about the origin or use of songs). Several participants also emphasized the importance of ensuring that the tool is optimized for mobile devices, as this would increase accessibility and convenience in various usage contexts. Users occasionally encountered technical issues, such as "score not found" errors or unclear explanations of node relationships within the phylogenetic tree. These problems arose because, as previously noted, the tree was not fully completed, and the nodes between genres were not functional in this project iteration.

Chapter 6

Conclusion

This chapter provides a critical assessment of the folk music visualization project and its contributions to computational ethnomusicology. **Section** 6.1 examines how well the implemented system fulfills its primary goal of making the diverse EA-DIGIFOLK corpus accessible and analyzable for ethnomusicologists, evaluating both technical achievements and the limitations identified through expert feedback. **Section** 6.2 outlines directions for future development, including the integration of additional musical traditions, enhanced interactivity, richer phylogenetic features, and audio playback features identified as priorities by both users and domain experts. Finally, **Section** 6.3 offers closing reflections on the project's broader significance, highlighting how the iterative, user-centered design process has advanced the field of music visualization and bridged traditional musicological practice with modern computational tools.

6.1 Main goals and alignment with the results

The outcomes of this research demonstrate that the primary objective, to develop a visualization tool capable of organizing and making accessible the extensive EA-DIGIFOLK music corpus, has been largely achieved.

Feedback from the co-design session and subsequent usability assessments confirms that the tool is both usable and effective, with an average SUS score of approximately 69.33 reflecting positive reception among participants. Users recognized the tool's value in enabling the exploration and comparison of large collections of scores, as well as its applicability in research and educational contexts.

At the same time, expert feedback identified clear directions for further improvement, including layout optimization, expanded features, and enhanced user guidance. These insights reinforce the importance of maintaining a user-centered, iterative approach to development, ensuring that future enhancements remain both practical and scalable.

By successfully integrating traditional musicological perspectives with modern computational techniques, this work not only advances the field of music information retrieval but also contributes meaningful tools and methodologies for the preservation and study of folk music. The strong alignment between the project's goals and the results achieved underscores the significance and potential impact of this research for both academic inquiry and professional practice in ethnomusicology.

6.2 Future Work

Although a satisfactory result was achieved in this project, there is still room for improvement in several aspects of the visualization tool. This section outlines potential enhancements that could not be implemented within the project's timeframe due to various constraints.

Additional Traditions

The main future implementation that we hope for this project is the incorporation of additional musical traditions, which would progressively enhance the tool's utility and scope. Having more traditions would open doors to even more comprehensive cross-cultural analysis and reveal even more patterns of musical evolution and influence across geographic regions.

The current implementation focuses on Irish and Galician traditions, providing a strong foundation for Celtic musical analysis. However, expanding to include other traditions such as Portuguese folk music (and others present on the EA-DIGIFOLK dataset) would enable researchers to explore wider connections and influences between different traditions.

This expansion would also benefit from collaborative partnerships with ethnomusicologists specializing in these various traditions, ensuring that the computational analysis respects and accurately represents the unique cultural contexts of each musical style. Such collaboration would further validate the tool's approach while making it increasingly valuable to a broader community of researchers.

Interactive Node Functionality

The most significant limitation in the current implementation is the restricted interactivity of internal nodes. While all leaf nodes (representing individual scores) are fully functional, the convergence points that represent shared musical patterns or evolutionary relationships could not be made interactive. These internal nodes would ideally provide valuable information about common structural elements between different musical pieces or genres. The complexity of properly mapping these convergence points to their respective set of songs, combined with technical implementation challenges, prevented this feature from being completed.

Enhanced Phylogenetic Tree Representation

One of the most consistent recommendations from both feedback sessions was the need to enhance the phylogenetic trees' ability to convey richer information. The current implementation, while functional, doesn't fully cover the visual encoding capabilities inherent to phylogenetic trees. In traditional phylogenetic analysis, visual properties serve crucial semantic roles: branch length represents evolutionary time or degree of divergence, while branch thickness indicates confidence levels or the strength of relationships.

Participants specifically highlighted how implementing these nuanced visual encodings would significantly enhance the tool's ability to communicate both the degree of similarity between musical pieces and the certainty of the proposed relationships. By incorporating branch length to represent musical distance and thickness to indicate statistical confidence in the relationships, users could intuitively grasp complex patterns of musical evolution at a glance.

This visual enhancement remains one of the highest-priority features for future development, as it would transform the tool into an even richer analytical instrument capable of revealing subtleties in musical relationships.

Integrated Audio Playback

Another valuable feature that couldn't be implemented was an integrated score player. Having the ability to hear each score directly within the application, both individually and in the comparison view, would create a more comprehensive analytical experience. This audio component would allow users to perceive similarities that might not be immediately apparent from visual inspection alone, creating a multi-sensory approach to musical analysis. The implementation would require integrating MIDI playback capabilities with the MusicXML rendering system already in place.

Performance Optimization

With larger datasets, the current implementation may face performance issues, particularly when rendering complex trees or multiple scores simultaneously. Implementing more efficient rendering techniques, lazy loading, and server-side processing for complex computations would improve the tool's scalability for larger musical corpora.

These potential improvements represent natural extensions of the current work and would address the feedback received during user testing. While they couldn't be implemented within the scope of this project, they provide a clear roadmap for future development to create an even more powerful tool for folk music analysis and comparison.

6.3 Conclusion

In conclusion, this research addresses the pressing challenge of organizing and interpreting the extensive and diverse datasets within the EA-DIGIFOLK project by developing a dedicated visualization tool. Through a user-centered, co-design approach, the project has demonstrated that effective visual representations can significantly enhance the accessibility and analytical potential of large musical corpora. The positive feedback from participants, reflected in both the SUS usability scores and qualitative insights, confirms that the tool aligns well with the needs of researchers and educators working with folk music data. By bridging traditional musicological practices with modern computational methods, this work not only contributes valuable tools and methodologies to the field of music information retrieval but also supports the ongoing preservation and deeper understanding of folk music traditions from Galicia, Ireland, and beyond.

References

Fur Elise (Ludwig Van Beethoven). URL http://archive.org/details/FurEliseLudwigVanBeethoven.

Shalbbya Ali, Safdar Tanweer, Syed Khalid, and Naseem Rao. Mel frequency cepstral coefficient: A review. 2021. doi: 10.4108/eai.27-2-2020.2303173.

Maciej Blaszke and Damian Koszewski. Determination of Low-Level Audio Descriptors of a Musical Instrument Sound Using Neural Network. In *2020 Signal Processing: Algorithms, Architectures, Arrangements, and Applications (SPA)*, pages 138–141, September 2020. doi: 10.23919/SPA50552.2020.9241264. URL https://ieeexplore.ieee.org/document/9241264/?arnumber=9241264. ISSN: 2326-0319.

John Brooke. Sus: A quick and dirty usability scale. *Usability Eval. Ind.*, 189, 11 1995.

Gabriel Dias Cantareira, Luis Gustavo Nonato, and Fernando V. Paulovich. MoshViz: A Detail+Overview Approach to Visualize Music Elements. *IEEE Transactions on Multimedia*, 18 (11):2238–2246, November 2016. ISSN 1520-9210, 1941-0077. doi: 10.1109/TMM.2016.2614226. URL http://ieeexplore.ieee.org/document/7577795/.

Nádia Carvalho, Sara Gonzalez-Gutierrez, Javier Merchan Sanchez-Jara, Gilberto Bernardes, and Maria Navarro-Cáceres. Encoding, Analysing and Modeling I-Folk: A New Database of Iberian Folk Music. In *8th International Conference on Digital Libraries for Musicology*, pages 75–83, Virtual Conference GA USA, July 2021. ACM. ISBN 978-1-4503-8429-2. doi: 10.1145/3469013.3469023. URL https://dl.acm.org/doi/10.1145/3469013.3469023.

Matthew Cooper, Jonathan Foote, Elias Pampalk, and George Tzanetakis. Visualization in Audio-Based Music Information Retrieval. *Computer Music Journal*, 30(2):42–62, June 2006. ISSN 0148-9267, 1531-5169. doi: 10.1162/comj.2006.30.2.42. URL https://direct.mit.edu/comj/article/30/2/42-62/94702.

Débora C. Corrêa and Francisco Ap. Rodrigues. A survey on symbolic data-based music genre classification. *Expert Systems with Applications*, 60:190–210, 2016. ISSN 0957-4174. doi: https://doi.org/10.1016/j.eswa.2016.04.008. URL https://www.sciencedirect.com/science/article/pii/S095741741630166X.

Roger T. Dean. *The Oxford Handbook of Computer Music*. Oxford Handbooks Series. Oxford University Press, Incorporated, Oxford, 1st ed edition, 2009. ISBN 978-0-19-533161-5 978-0-19-971593-0.

Pierre Dragicevic. 3D Spectrogram, June 2018. URL https://dataphys.org/list/3d-spectrogram/. Section: entries.

Steven Feld and Bruno Nettl. The Study of Ethnomusicology. *Latin American Music Review / Revista de Música Latinoamericana*, 7(2):375, 1986. ISSN 01630350. doi: 10.2307/780225. URL https://www.jstor.org/stable/780225?origin=crossref.

Jonathan Foote. Visualizing music and audio using self-similarity. In *Proceedings of the seventh ACM international conference on Multimedia (Part 1)*, pages 77–80, Orlando Florida USA, October 1999. ACM. ISBN 978-1-58113-151-2. doi: 10.1145/319463.319472. URL https://dl.acm.org/doi/10.1145/319463.319472.

Liu Guangchao, Lee Yok Fee, Ratna Roshida Ab. Razak, and Arfah Ab. Majid. Folk Song Conservation Strategies from a Cross-Cultural Perspective: A Systematic Literature Review. *International Journal of Academic Research in Business and Social Sciences*, 14(6):Pages 1250–1272, June 2024. ISSN 2222-6990. doi: 10.6007/IJARBSS/v14-i6/21785. URL https://hrmars.com/journals/papers/IJARBSS/v14-i6/21785.

Dilek Göktürk. Kodály and Orff: A comparison of two approaches in early music education. 8 (15), 2012.

L. J. Heyer, S. Kruglyak, and S. Yooseph. Exploring expression data: identification and analysis of coexpressed genes. *Genome Research*, 9(11):1106–1115, November 1999. ISSN 1088-9051. doi: 10.1101/gr.9.11.1106.

Jasmine Irani, Nitin Pise, and Madhura Phatak. Clustering Techniques and the Similarity Measures used in Clustering: A Survey. *International Journal of Computer Applications*, 134(7):9–14, January 2016. ISSN 09758887. doi: 10.5120/ijca2016907841. URL http://www.ijcaonline.org/research/volume134/number7/irani-2016-ijca-907841.pdf.

Eric Isaacson. WHAT YOU SEE IS WHAT YOU GET: ON VISUALIZING MUSIC. 2005.

R. Khulusi, J. Kusnick, C. Meinecke, C. Gillmann, J. Focht, and S. Jänicke. A Survey on Visualizations for Musical Data. *Computer Graphics Forum*, 39(6):82–110, September 2020. ISSN 0167-7055, 1467-8659. doi: 10.1111/cgf.13905. URL https://onlinelibrary.wiley.com/doi/10.1111/cgf.13905.

Andy kolovos. Archiving culture: American folklore archives in theory and practice. 2010.

Agata Krajewska. Digitisation of Folklore Archives: A Crisis of Tradition or Its 'New Life' on the Internet? The Example of Adolf Dygacz's Collection. *Kwartalnik Młodych Muzykologów UJ*, (46 (3)):67–86, 2020. ISSN 23537094. doi: 10.4467/23537094KMMUJ.20.038.13911. URL https://ejournals.eu/en/journal/kmmuj/article/digitisation-of-folklore-archives-a-crisis-of-tradition-or-its-new-life-on-th

Alexander Lerch. *An introduction to audio content analysis: music information retrieval tasks & applications*. John Wiley & Sons, Inc, Hoboken, New Jersey, second edition edition, 2023. ISBN 978-1-119-89096-6 978-1-119-89097-3 978-1-119-89098-0.

Elad Liebman, Eitan Ornoy, and Benny Chor. A Phylogenetic Approach to Music Performance Analysis. *Journal of New Music Research*, 41(2):195–222, June 2012. ISSN 0929-8215, 1744-5027. doi: 10.1080/09298215.2012.668194. URL http://www.tandfonline.com/doi/abs/10.1080/09298215.2012.668194.

Hugo B. Lima, Carlos G. R. Dos Santos, and Bianchi S. Meiguins. A Survey of Music Visualization Techniques. *ACM Computing Surveys*, 54(7):1–29, September 2022. ISSN 0360-0300, 1557-7341. doi: 10.1145/3461835. URL https://dl.acm.org/doi/10.1145/3461835.

Lele Liu, Veronica Morfi, and Emmanouil Benetos. Joint Multi-Pitch Detection and Score Transcription for Polyphonic Piano Music. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 281–285, Toronto, ON, Canada, June 2021. IEEE. ISBN 978-1-7281-7605-5. doi: 10.1109/ICASSP39728.2021.9413601. URL https://ieeexplore.ieee.org/document/9413601/.

Wilson Louie. MusicPlot: Interactive Self-Similarity Matrix for Music Structure Visualization. 2020.

Delfina Malandrino, Donato Pirozzi, Gianluca Zaccagnino, and Rocco Zaccagnino. A Color-Based Visualization Approach to Understand Harmonic Structures of Musical Compositions. In *2015 19th International Conference on Information Visualisation*, pages 56–61, July 2015. doi: 10.1109/iV.2015.21. URL https://ieeexplore.ieee.org/document/7272579/?arnumber=7272579. ISSN: 2375-0138.

Samuel D. Miller. Guido D'Arezzo: Medieval Musician and Educator. *Journal of Research in Music Education*, 21(3):239–245, October 1973. ISSN 0022-4294, 1945-0095. doi: 10.2307/3345093. URL https://journals.sagepub.com/doi/10.2307/3345093.

Meinard Müller. *Fundamentals of Music Processing: Using Python and Jupyter Notebooks*. Springer International Publishing, Cham, 2021. ISBN 978-3-030-69807-2 978-3-030-69808-9. doi: 10.1007/978-3-030-69808-9. URL https://link.springer.com/10.1007/978-3-030-69808-9.

Saul B. Needleman and Christian D. Wunsch. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *Journal of Molecular Biology*, 48(3): 443–453, March 1970. ISSN 0022-2836. doi: 10.1016/0022-2836(70)90057-4. URL https://www.sciencedirect.com/science/article/pii/0022283670900574.

Bruno Nettl. Folk and Traditional Music of the Western Continents. 1965.

Bruno Nettl. Folk music | Definition, History, Artists, Songs, Bands, Instruments, & Facts | Britannica, January 2025. URL https://www.britannica.com/art/folk-music.

Oriol Nieto, Gautham J. Mysore, Cheng-i Wang, Jordan B. L. Smith, Jan Schlüter, Thomas Grill, and Brian McFee. Audio-Based Music Structure Analysis: Current Trends, Open Challenges, and Applications. *Transactions of the International Society for Music Information Retrieval*, 3(1):246–263, December 2020. ISSN 2514-3298. doi: 10.5334/tismir.54. URL http://transactions.ismir.net/articles/10.5334/tismir.54/.

Maria Panteli, Emmanouil Benetos, and Simon Dixon. A review of manual and computational approaches for the study of world music corpora. *Journal of New Music Research*, 47(2):176–189, March 2018. ISSN 0929-8215, 1744-5027. doi: 10.1080/09298215.2017.1418896. URL https://www.tandfonline.com/doi/full/10.1080/09298215.2017.1418896.

Hilda Romero-Velo, Gilberto Bernardes, Susana Ladra, Jos'e R. Param'a, and Fernando Silva-Coira. Phylo-analysis of folk traditions: A methodology for the hierarchical musical similarity analysis. In *Proceedings of the 26th International Society for Music Information Retrieval Conference (ISMIR)*, Daejeon, South Korea, 2025. ISMIR.

Eric Thul and Godfried T Toussaint. A Comparative Phylogenetic-Tree Analysis of African Timelines and North Indian Talas. 2008.

M. Wattenberg. Arc diagrams: visualizing structure in strings. In *IEEE Symposium on Information Visualization, 2002. INFOVIS 2002.*, pages 110–116, Boston, MA, USA, 2002. IEEE Comput. Soc. ISBN 978-0-7695-1751-3. doi: 10.1109/INFVIS.2002.1173155. URL http://ieeexplore.ieee.org/document/1173155/.

Hendry Wijaya and Raymond Sunardi Oetama. Song Similarity Analysis With Clustering Method On Korean Pop Song. In *2021 6th International Conference on New Media Studies (CONMEDIA)*, pages 66–71, Tangerang, Indonesia, October 2021. IEEE. ISBN 978-1-6654-3353-2 978-1-6654-3354-9. doi: 10.1109/CONMEDIA53104.2021.9617204. URL https://ieeexplore.ieee.org/document/9617204/.

Tiange Zhu, Raphaël Fournier-S'niehotta, Philippe Rigaux, and Nicolas Travers. A Framework for Content-Based Search in Large Music Collections. *Big Data and Cognitive Computing*, 6(1):23, March 2022. ISSN 2504-2289. doi: 10.3390/bdcc6010023. URL https://www.mdpi.com/2504-2289/6/1/23. Number: 1 Publisher: Multidisciplinary Digital Publishing Institute.

Appendix A

Co-Design Session Script

A.1 Co-Design Session

First, I would like to ask you if I can record this session for later analysis and eventual transcription. It will be used for research purposes only, and it will be destroyed once the project is finished. Hello everyone, my name is Pedro Magalhães, and I'm currently enrolled in the Multimedia Master's at the University of Porto. I'm writing a thesis and developing a project on visualizing structural relationships in Irish and Galician folk music, which can hopefully be extended to other traditions in the future. I'm grateful to have you here to run this co-design session. Purpose of the Session: The goal of this session is to gather your feedback on an interactive visualization I'm developing. The visualization includes three main components:

1. Phylogenetic trees exploring structural relationships between genres within each tradition.

2. A projection of the dataset as spatial points in a map to navigate clusters' resemblance based on chromatic, rhythmic, and a combination of these characteristics.

3. Score visualizations with enhanced features for comparing different tunes and songs by exposing similar patterns.

This study is part of the European project EA-DIGIFOLK, which you are already familiar with. Since this tool is meant to be used by scholars and practitioners, like yourselves, I focused on designing a system that helps streamline exploration, interpretation, and comparison, making it as intuitive and flexible as possible, but it is mostly directed to experts.

A.2 Session Format

We'll be working with some mockups that I designed in a software called Figma. Don't worry if you've never used Figma, it's designed to be very intuitive. Also, I will walk you through the prototypes and ask you to do specific tasks as we proceed. During these tasks, I would like to hear your thoughts on the visualization and in which ways I could improve it. During the Session, I'd love to hear your thoughts on:

1. Readability and interpretation
   - Are the visual representations clear?
   - Is the use of color, labels, or grouping helpful or confusing?

2. Filter parametrization
   - What kind of filters would you expect to use?

3. Design and usability
   - Is the navigation intuitive?
   - Is it clear what each part of the interface is for?

4. Research expectations
   - Would this kind of tool be helpful in your research?
   - What kind of tasks would you want to accomplish here?
   - Is there anything missing or unclear?

5. Additional suggestions
   - How would you imagine using this in a real analysis or publication?
   - What would make this tool more powerful or adaptable to your needs?

Please feel free to speak aloud as you go — your thought process is helpful. I might occasionally ask you to elaborate or reflect on something you've said. I'll also take notes throughout. At the end, I'll open the floor for any final thoughts, comments, or wishlist items. The experiment will be segmented into 4 paths, which we will follow through the mockups.

Appendix B

Co-Design Session

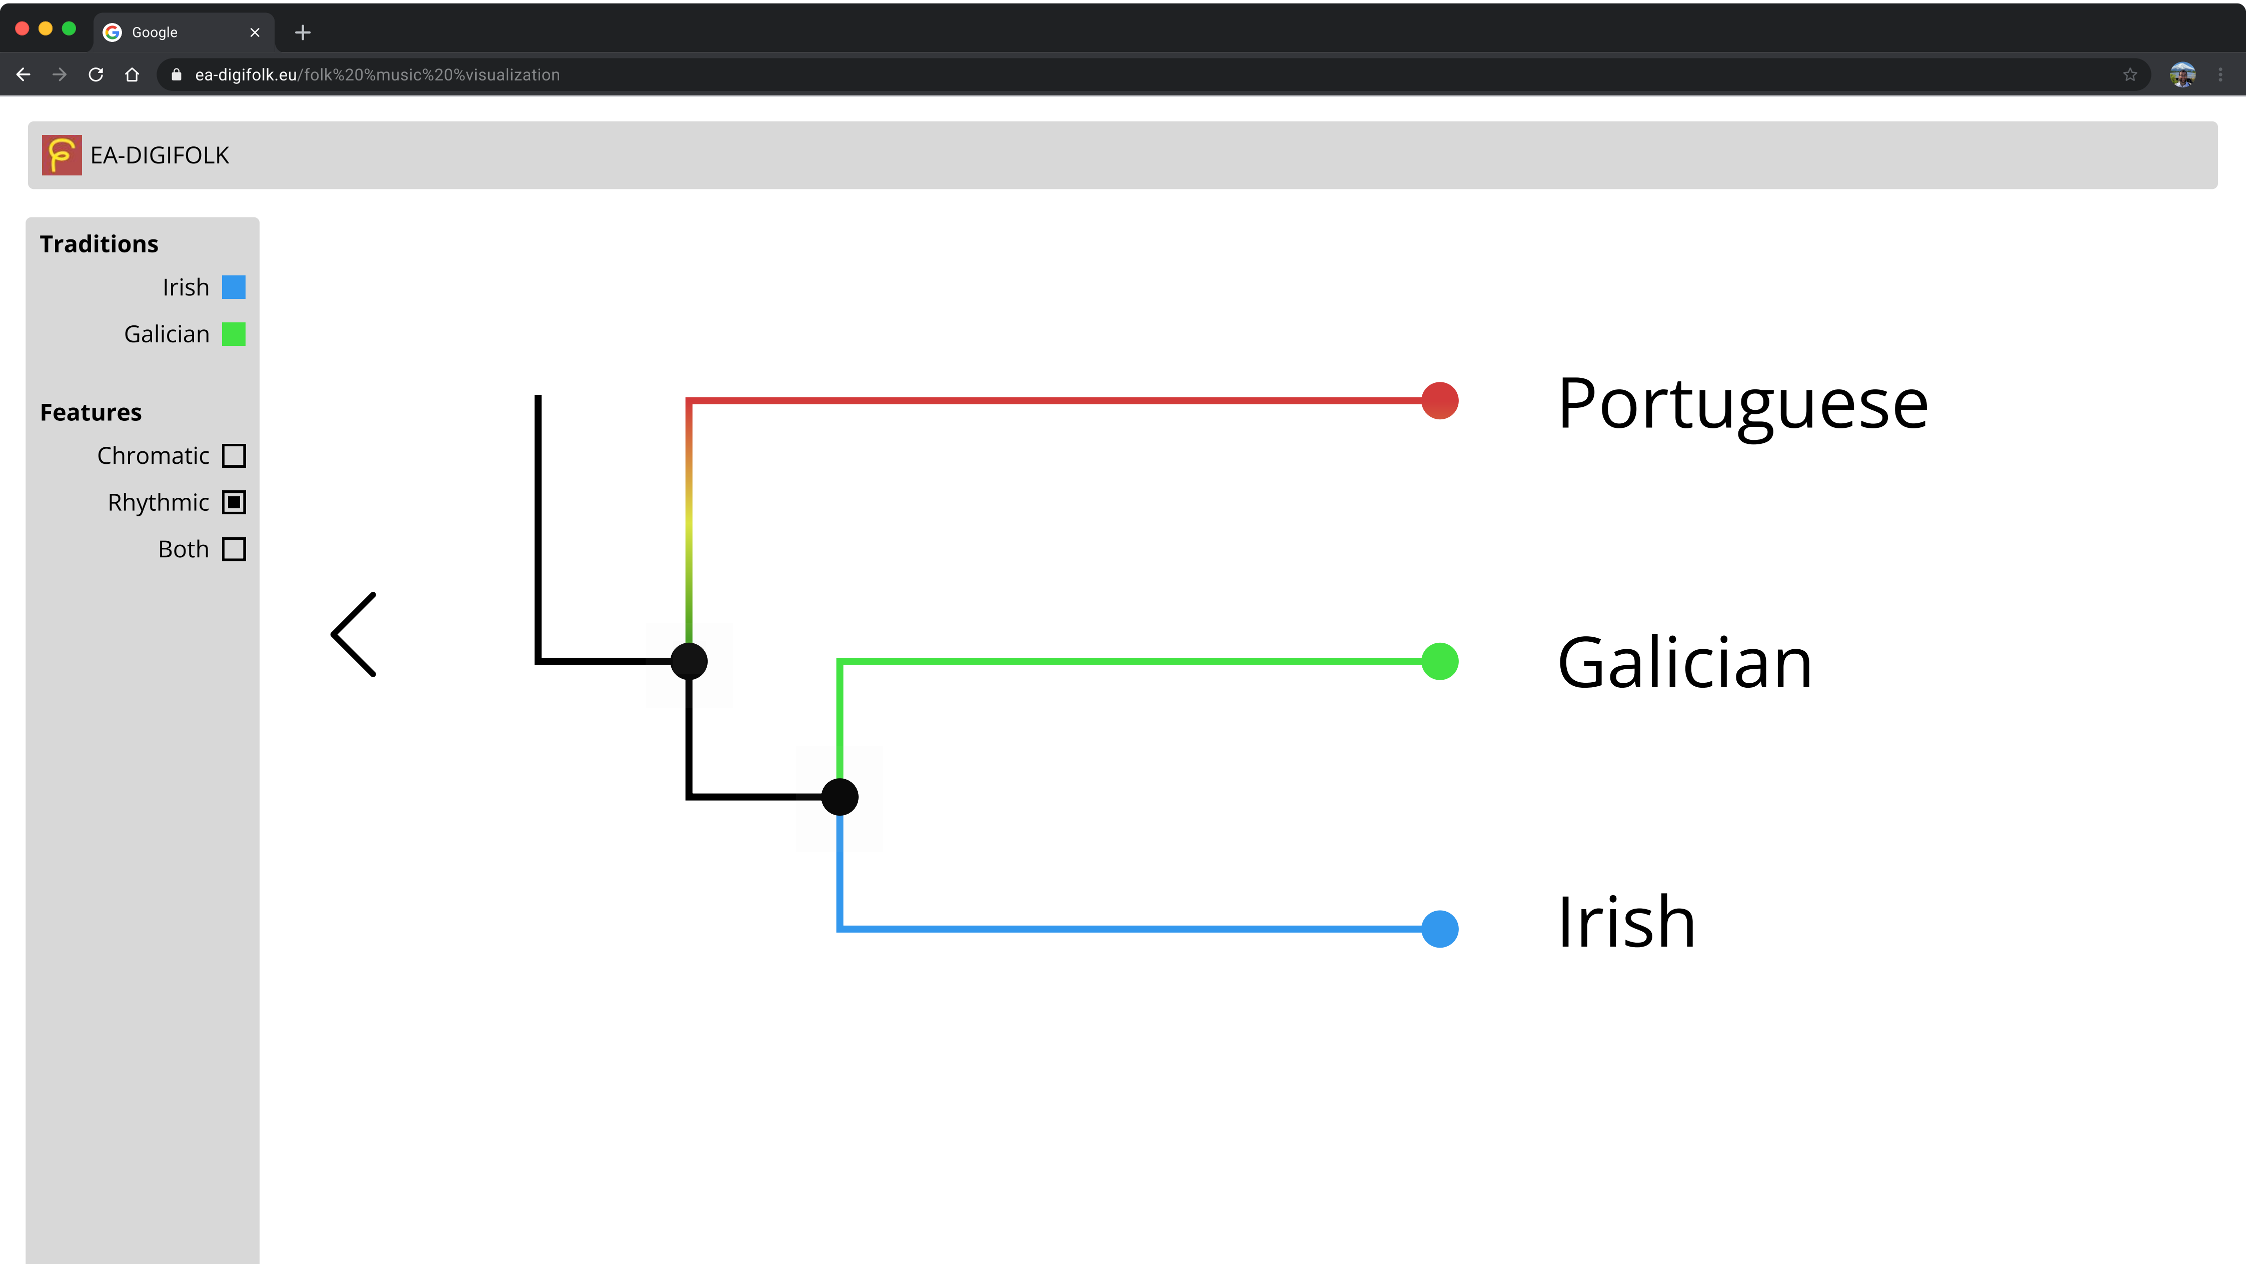
Google
ea-digifolk.eu/folk%20%music%20%visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
Portuguese
Galician
Irish


Google
ea-digifolk.eu/folk%20%music%20%visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
alalas
polka
slide
mazurcas
marchas
rumbas
polca
waltz
pasodobles
foliadas
jotas
valses
pasacorredoi
muineiras
jig
mazurka
march
hornpipe
barndance
reel
strathspey


Google
ea-digifolk.eu/folk%20music%20visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
muineiras
pasacorredoi
mazurka
jig
hornpipe
march
barndance
reel
strathspey
alalas
slide
polka
rumbas
polca
waltz
pasodobles
valses
foliada
jotas
marchas
mazurcas

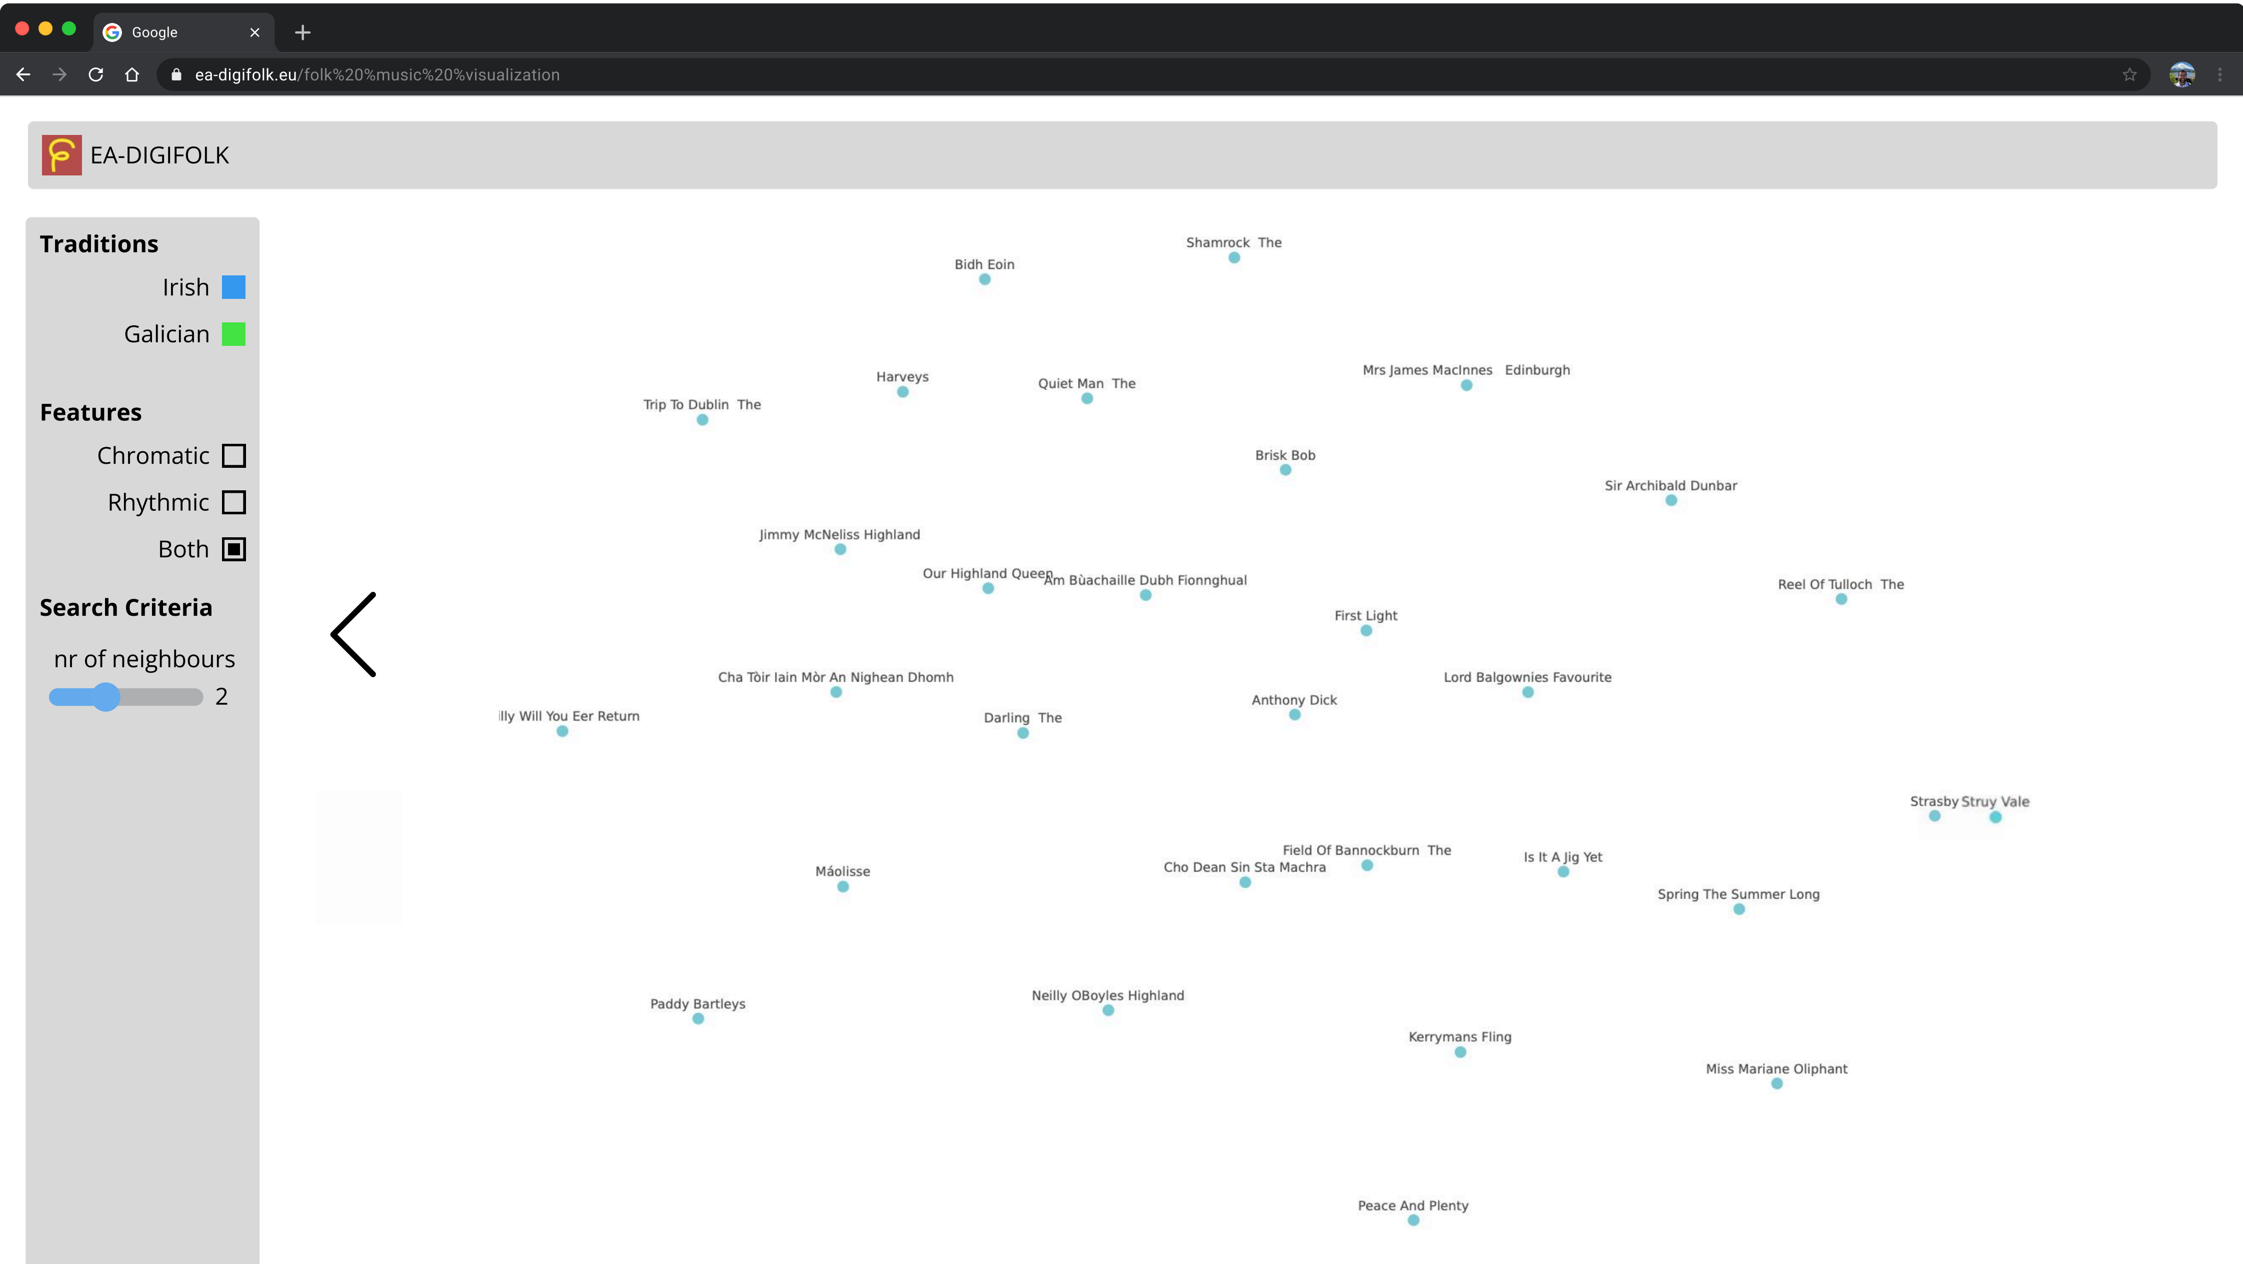
Google
ea-digifolk.eu/folk%20music%20visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
Search Criteria
nr of neighbours
2
Shamrock The
Bidh Eoin
Harveys
Quiet Man The
Mrs James MacInnes Edinburgh
Trip To Dublin The
Brisk Bob
Sir Archibald Dunbar
Jimmy McNeliss Highland
Our Highland Queen
Am Bùachaille Dubh Fionnghual
Reel Of Tulloch The
First Light
Cha Tòir Iain Mòr An Nighean Dhomh
Lord Balgownies Favourite
Anthony Dick
illy Will You Eer Return
Darling The
Strasby Struy Vale
Field Of Bannockburn The
Is It A Jig Yet
Cho Dean Sin Sta Machra
Máolisse
Spring The Summer Long
Paddy Bartleys
Neilly OBoyles Highland
Kerrymans Fling
Miss Mariane Oliphant
Peace And Plenty

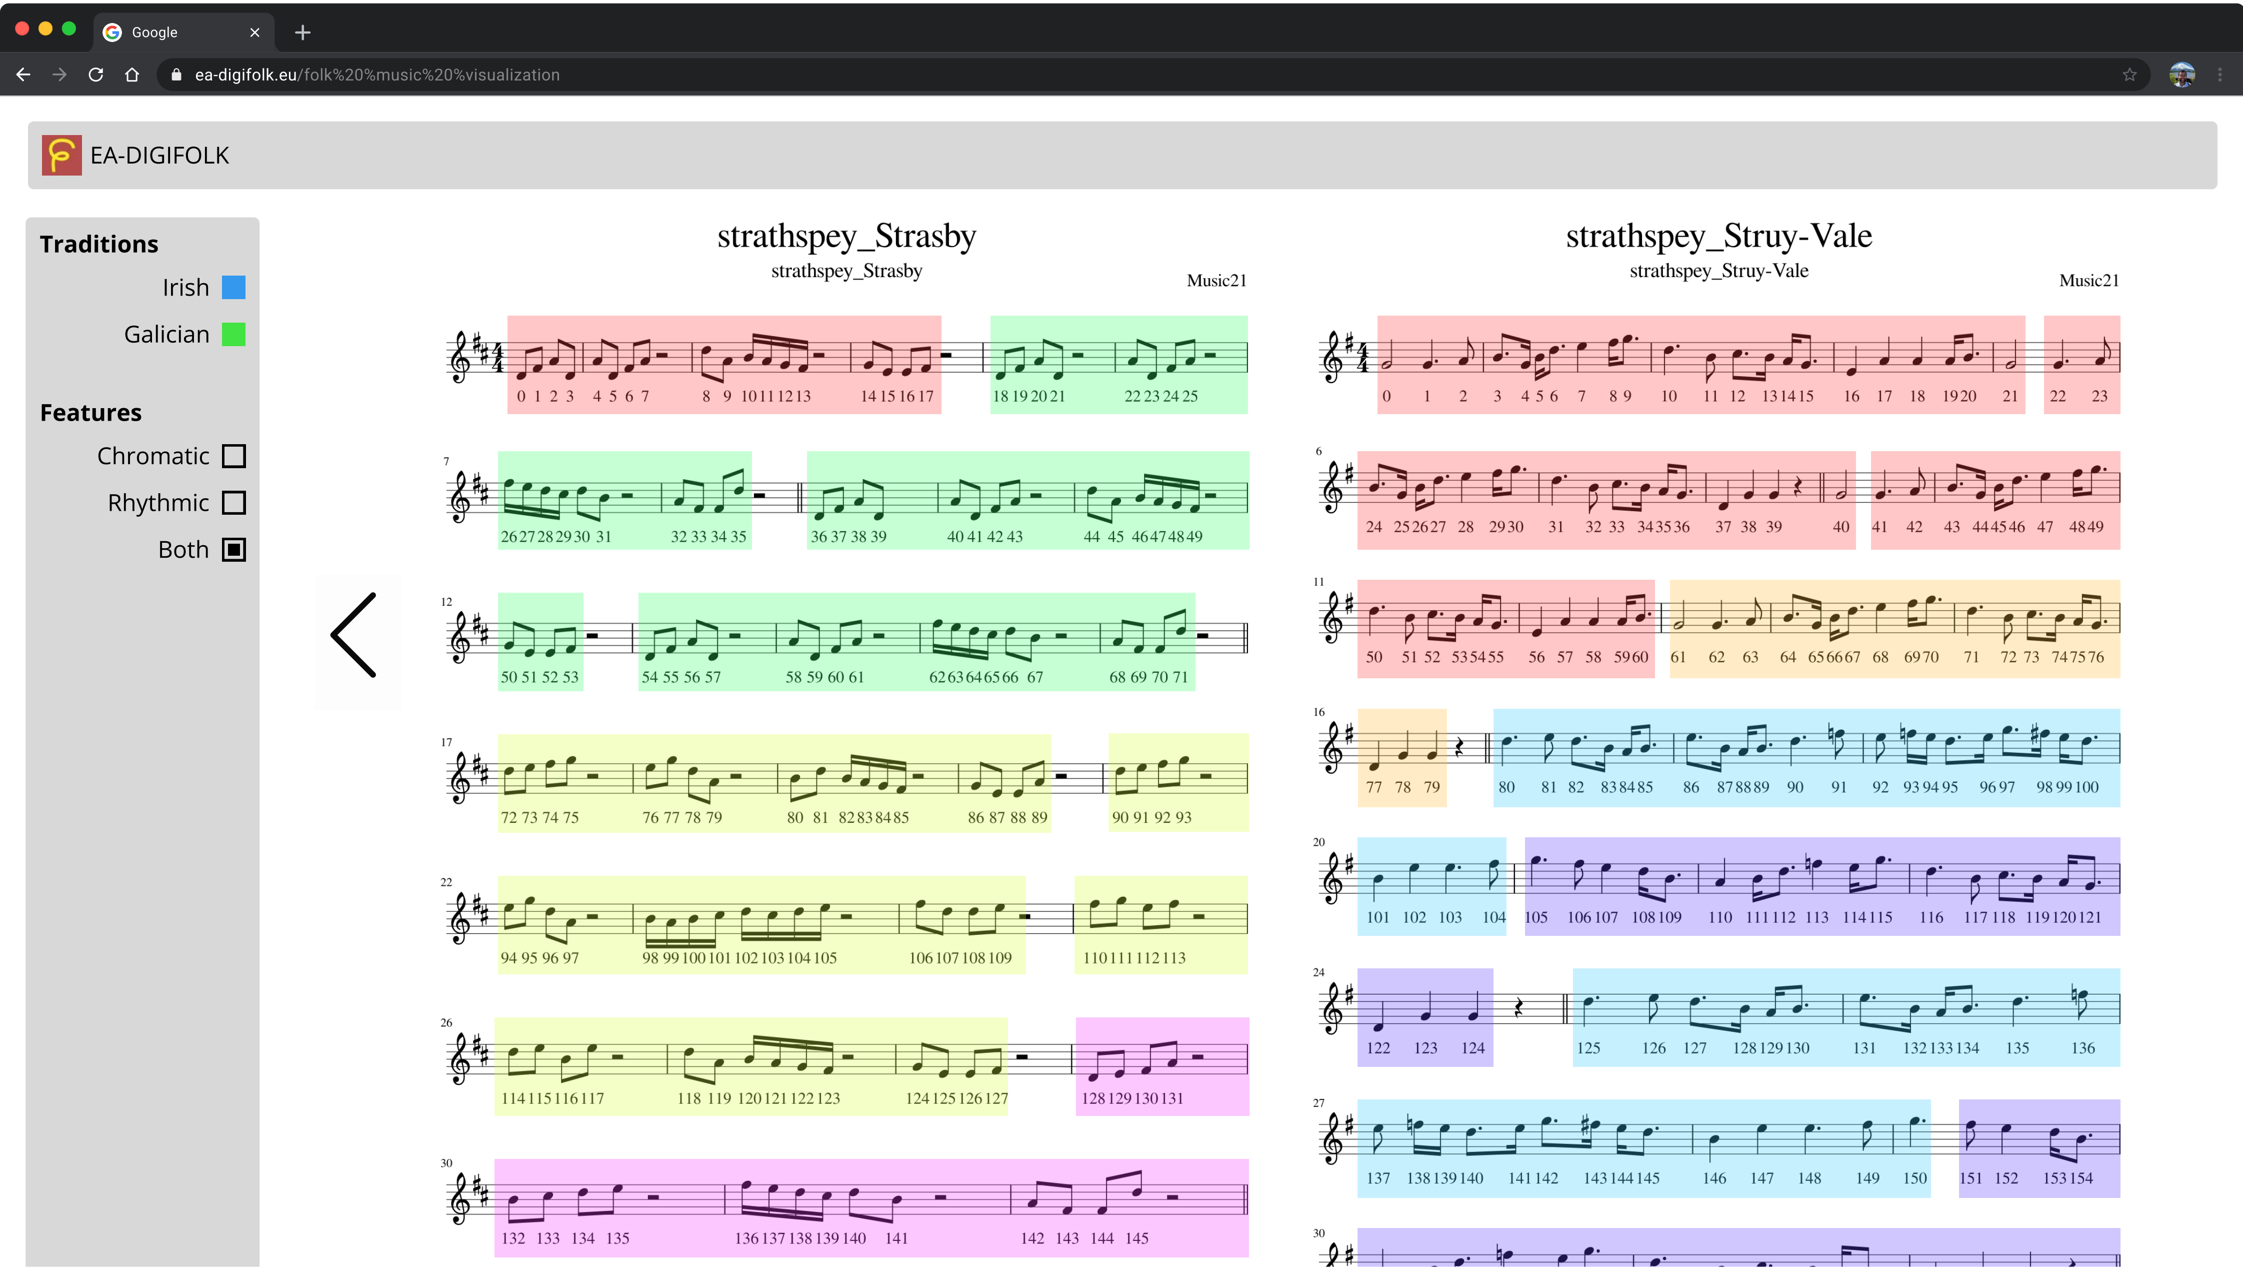
Google
ea-digifolk.eu/folk%20music%20visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
strathspey_Strasby
strathspey_Strasby
Music21
strathspey_Struy-Vale
strathspey_Struy-Vale
Music21

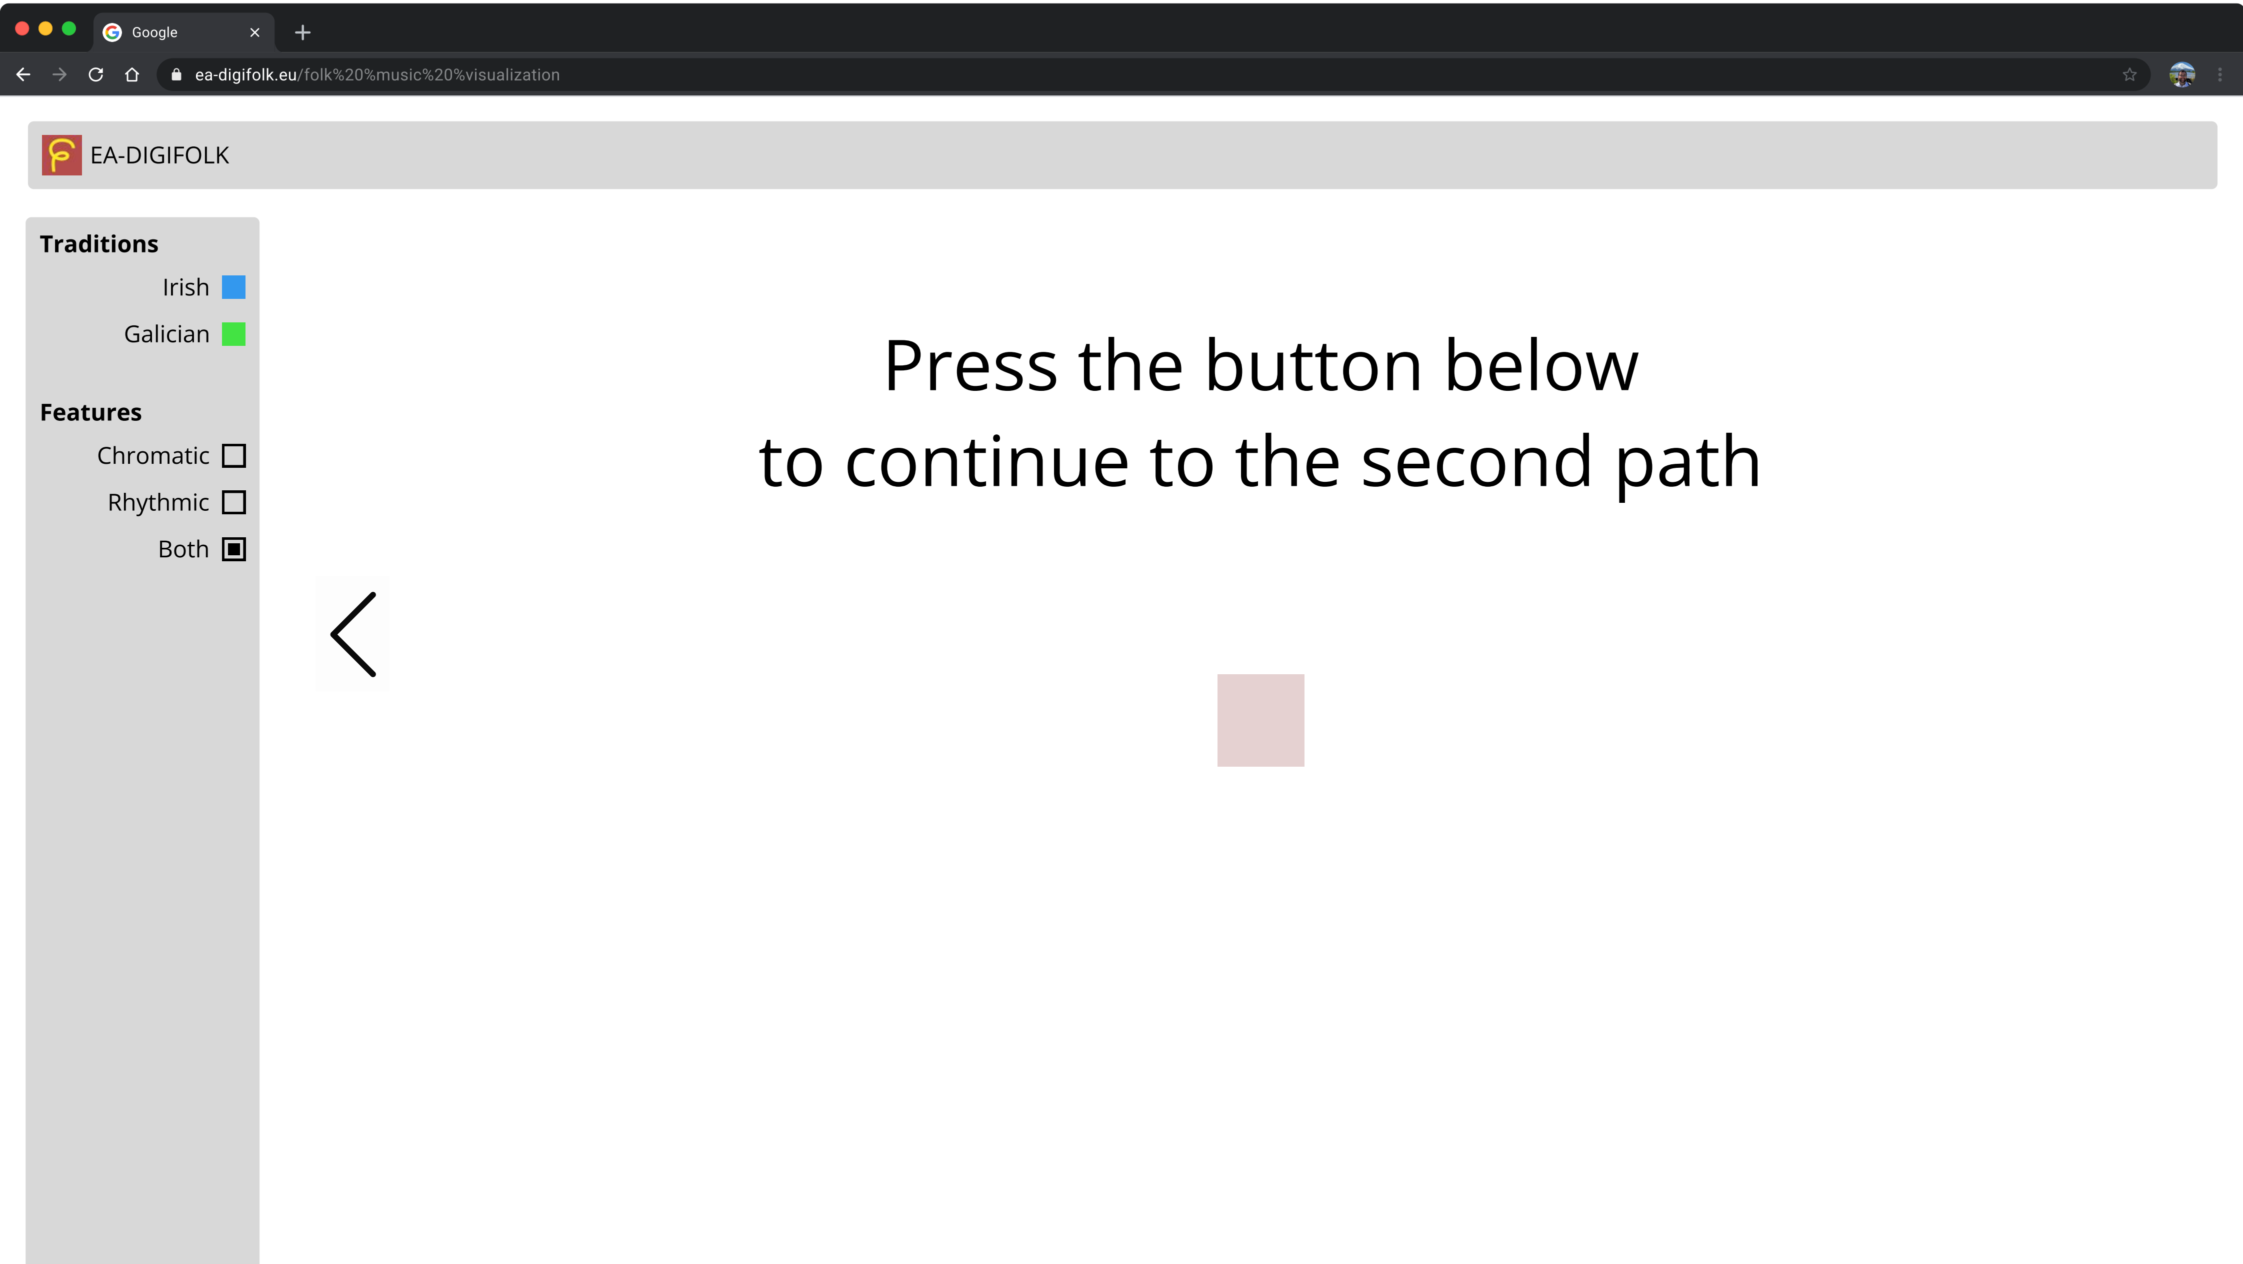
Google
ea-digifolk.eu/folk%20%music%20%visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
Press the button below
to continue to the second path


Google
ea-digifolk.eu/folk%20%music%20%visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
alalas
polka
slide
mazurcas
marchas
rumbas
polca
waltz
pasodobles
foliadas
jotas
valses
pasacorredoi
muineiras
jig
mazurka
march
hornpipe
barndance
reel
strathspey

Google
ea-digifolk.eu/folk%20music%20visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
muineiras
pasacorredoi
mazurka
jig
hornpipe
march
barndance
reel
strathspey
alalas
slide
polka
rumbas
polca
waltz
pasodobles
valses
foliada
jotas
marchas
mazurcas

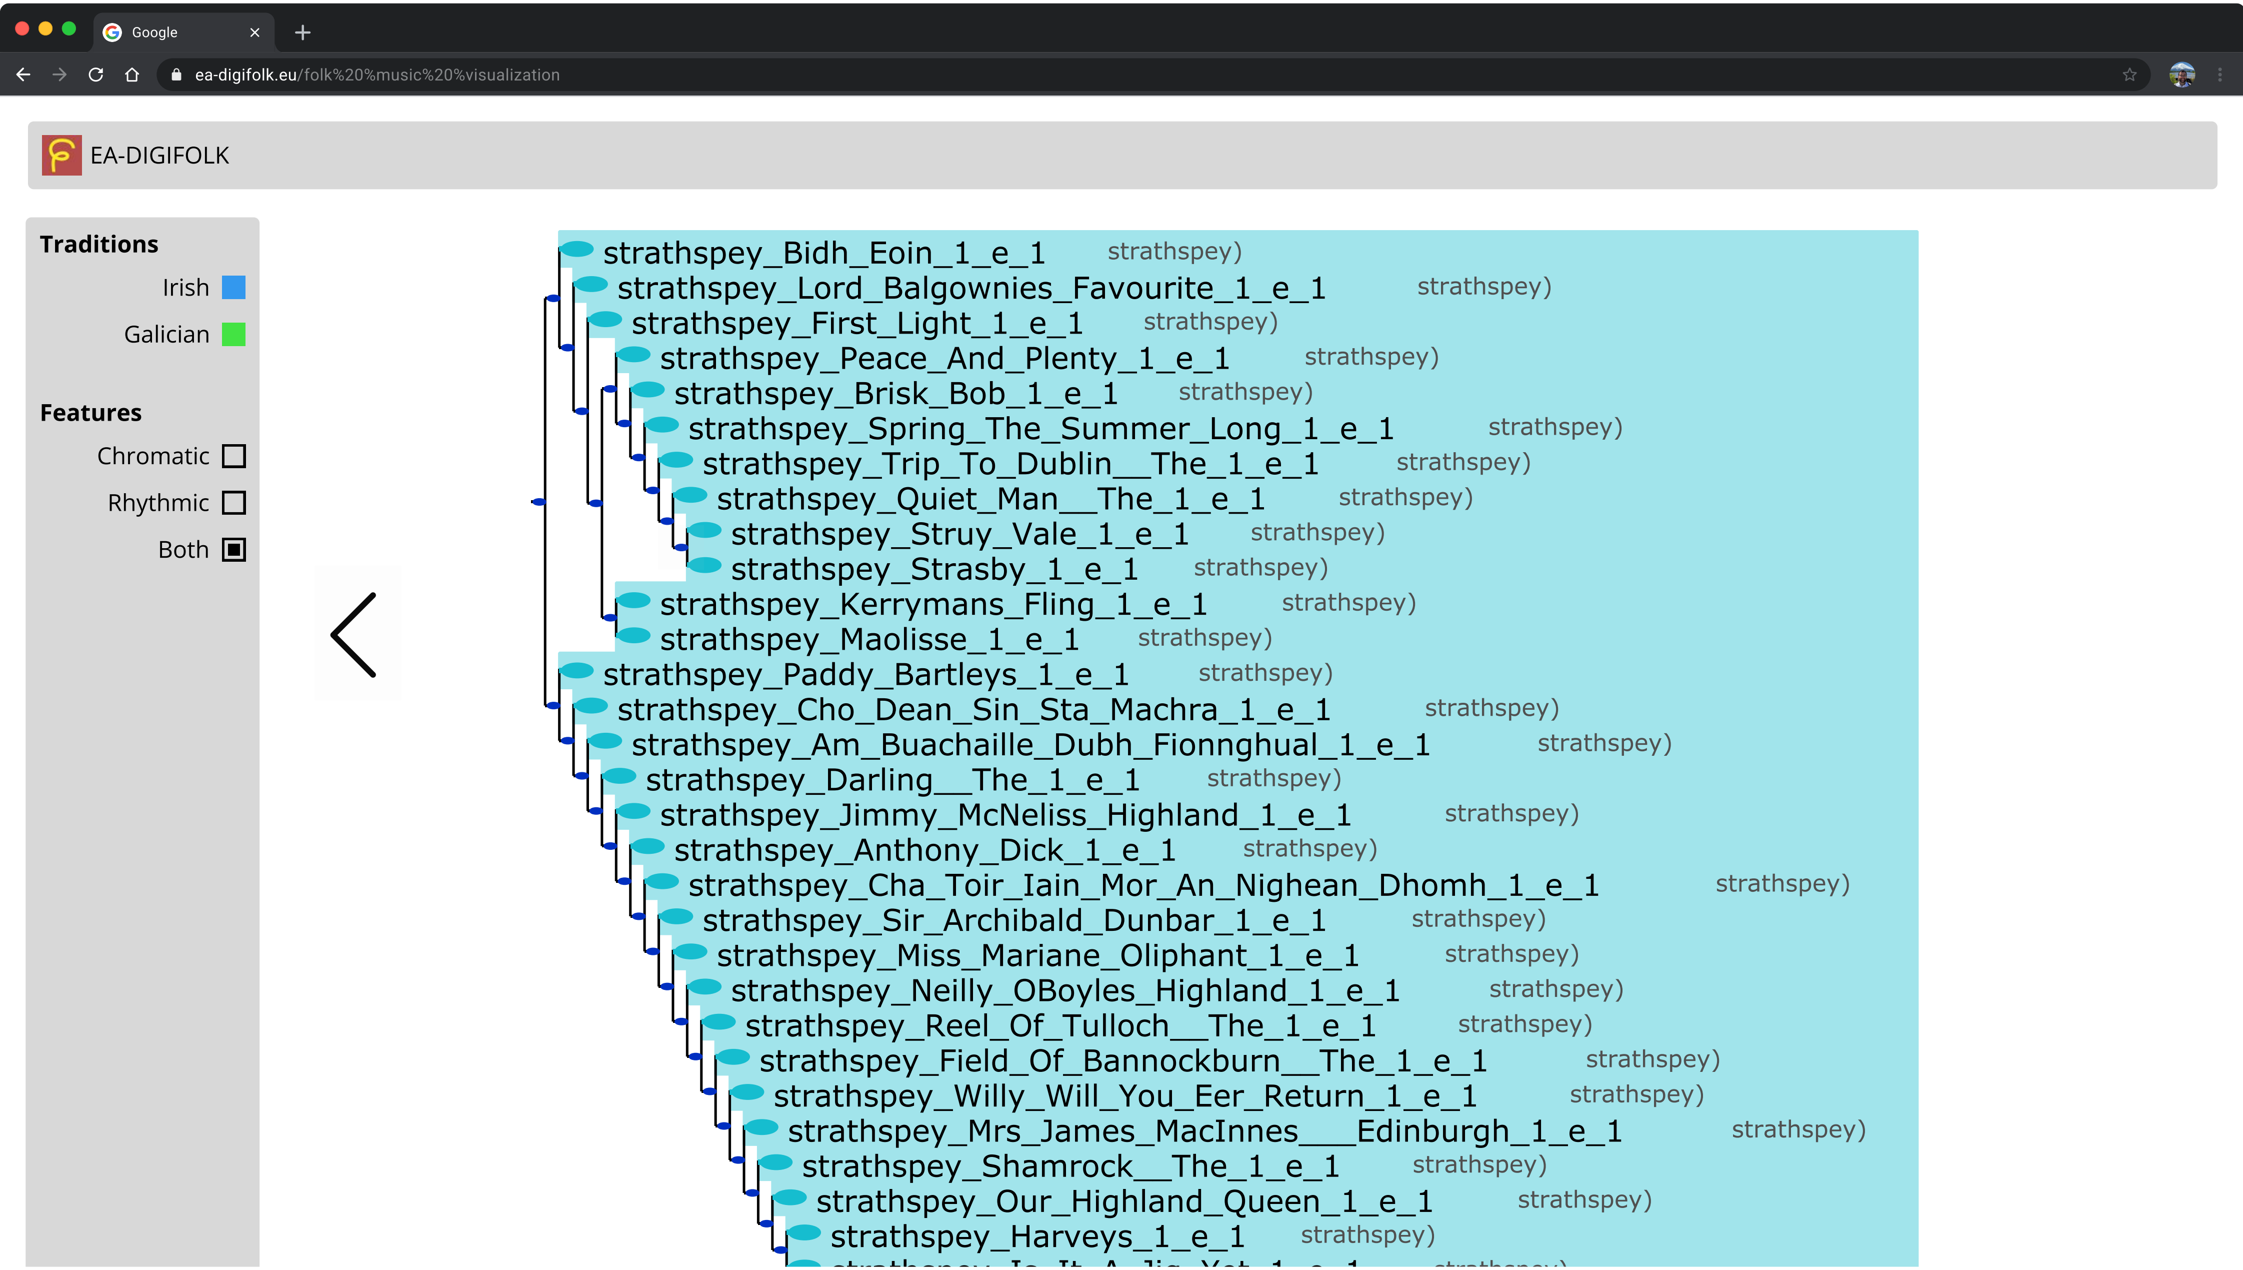
Google
ea-digifolk.eu/folk%20music%20visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
strathspey_Bidh_Eoin_1_e_1 strathspey)
strathspey_Lord_Balgownies_Favourite_1_e_1 strathspey)
strathspey_First_Light_1_e_1 strathspey)
strathspey_Peace_And_Plenty_1_e_1 strathspey)
strathspey_Brisk_Bob_1_e_1 strathspey)
strathspey_Spring_The_Summer_Long_1_e_1 strathspey)
strathspey_Trip_To_Dublin__The_1_e_1 strathspey)
strathspey_Quiet_Man__The_1_e_1 strathspey)
strathspey_Struy_Vale_1_e_1 strathspey)
strathspey_Strasby_1_e_1 strathspey)
strathspey_Kerrymans_Fling_1_e_1 strathspey)
strathspey_Maolisse_1_e_1 strathspey)
strathspey_Paddy_Bartleys_1_e_1 strathspey)
strathspey_Cho_Dean_Sin_Sta_Machra_1_e_1 strathspey)
strathspey_Am_Buachaille_Dubh_Fionnghual_1_e_1 strathspey)
strathspey_Darling__The_1_e_1 strathspey)
strathspey_Jimmy_McNeliss_Highland_1_e_1 strathspey)
strathspey_Anthony_Dick_1_e_1 strathspey)
strathspey_Cha_Toir_Iain_Mor_An_Nighean_Dhomh_1_e_1 strathspey)
strathspey_Sir_Archibald_Dunbar_1_e_1 strathspey)
strathspey_Miss_Mariane_Oliphant_1_e_1 strathspey)
strathspey_Neilly_OBoyles_Highland_1_e_1 strathspey)
strathspey_Reel_Of_Tulloch__The_1_e_1 strathspey)
strathspey_Field_Of_Bannockburn__The_1_e_1 strathspey)
strathspey_Willy_Will_You_Eer_Return_1_e_1 strathspey)
strathspey_Mrs_James_MacInnes___Edinburgh_1_e_1 strathspey)
strathspey_Shamrock__The_1_e_1 strathspey)
strathspey_Our_Highland_Queen_1_e_1 strathspey)
strathspey_Harveys_1_e_1 strathspey)

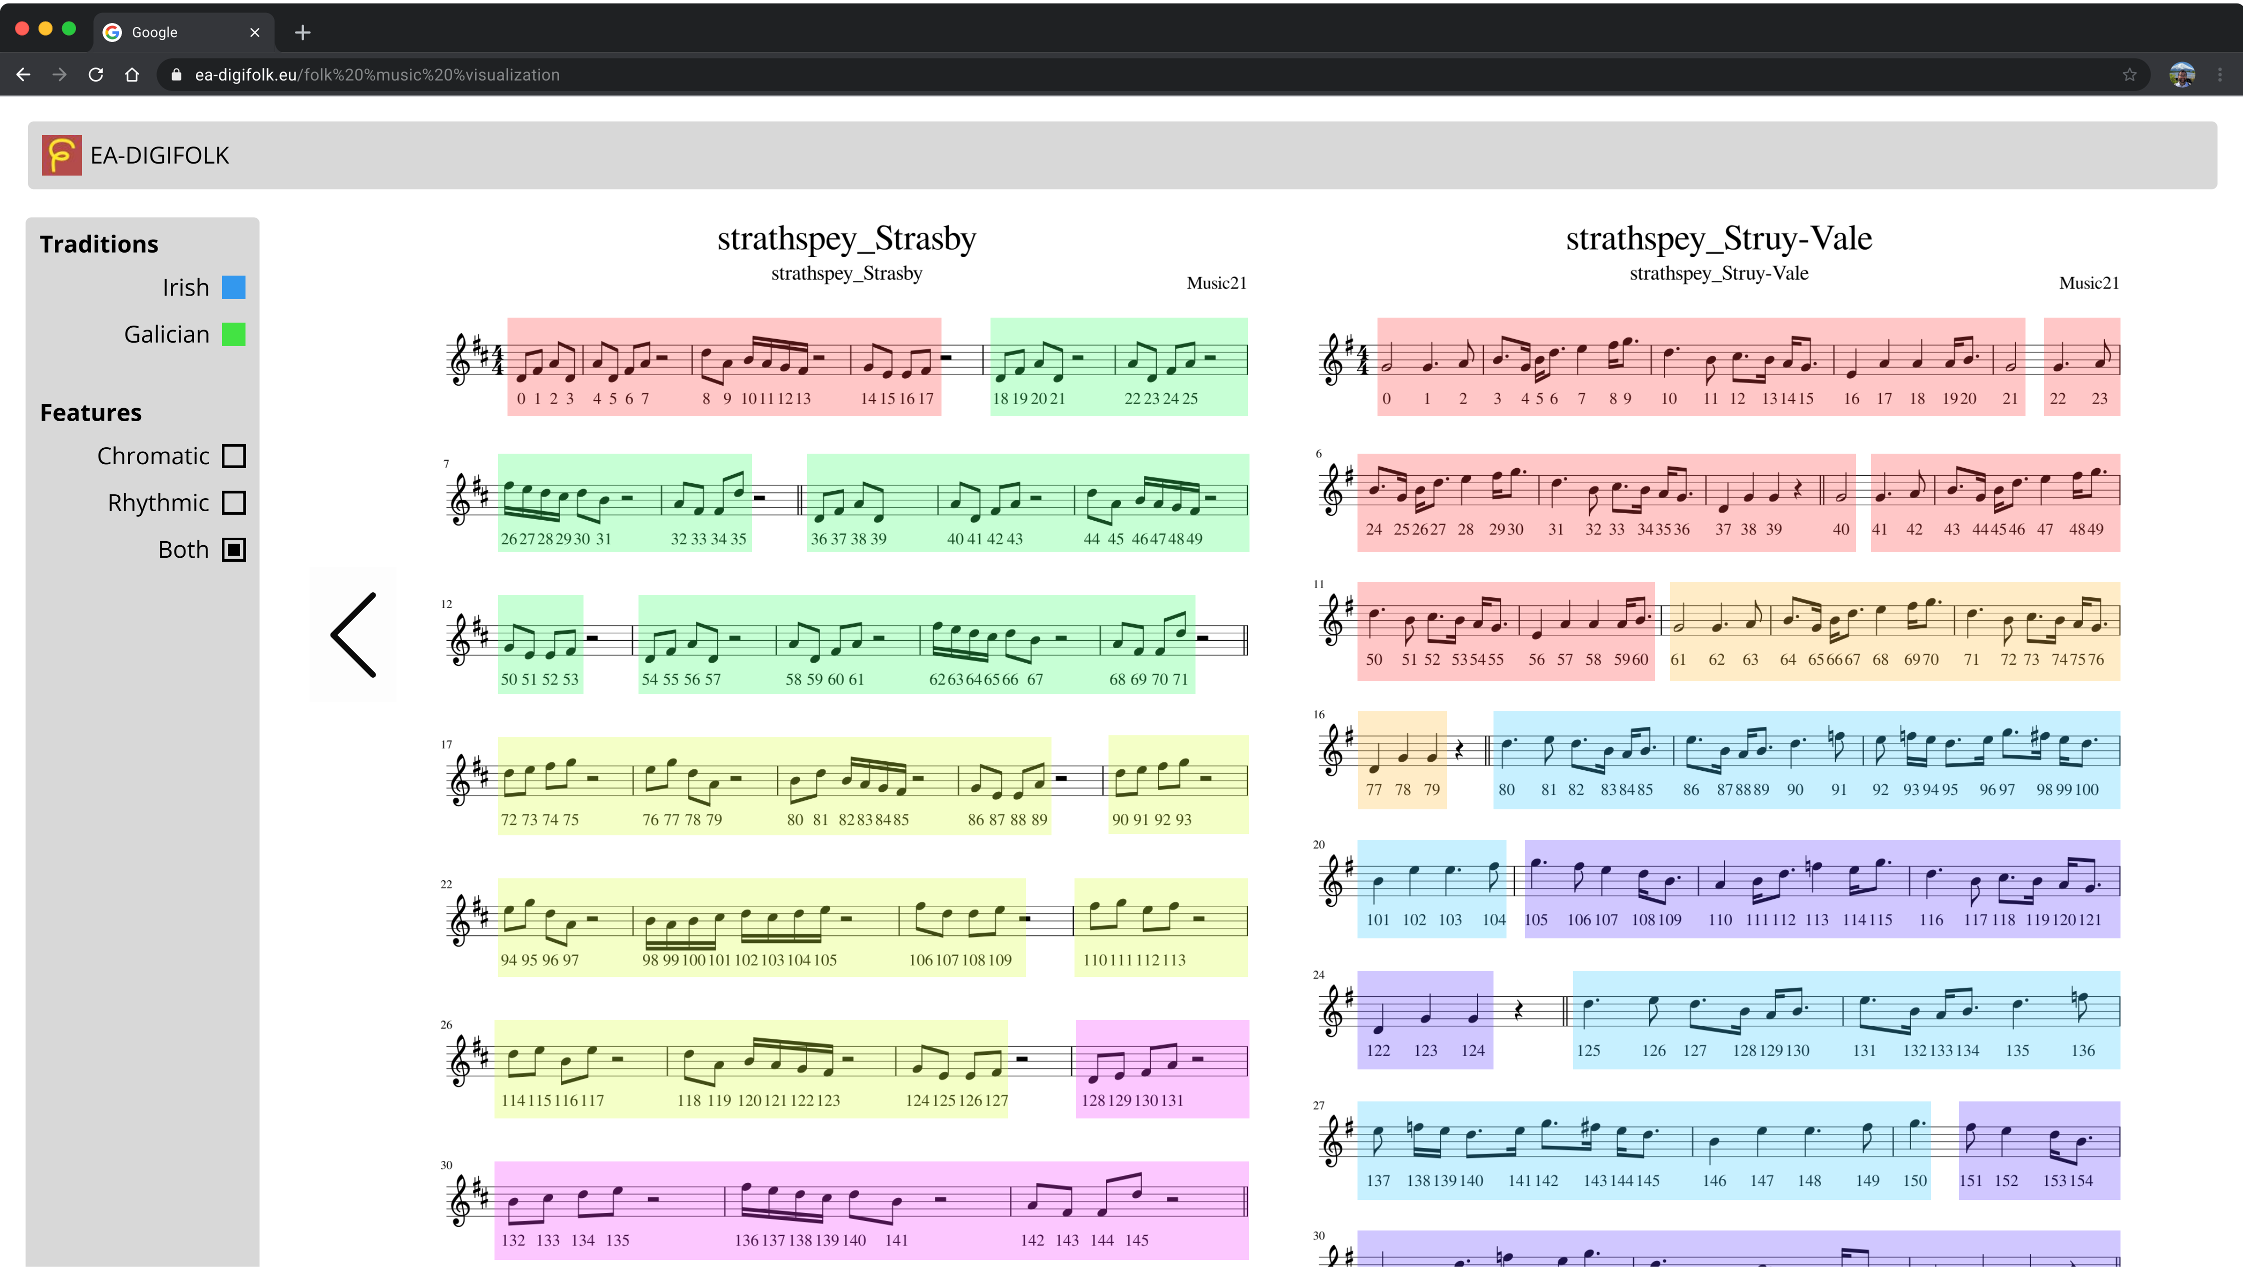
Google
ea-digifolk.eu/folk%20music%20visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
strathspey_Strasby
strathspey_Strasby
Music21
strathspey_Struy-Vale
strathspey_Struy-Vale
Music21

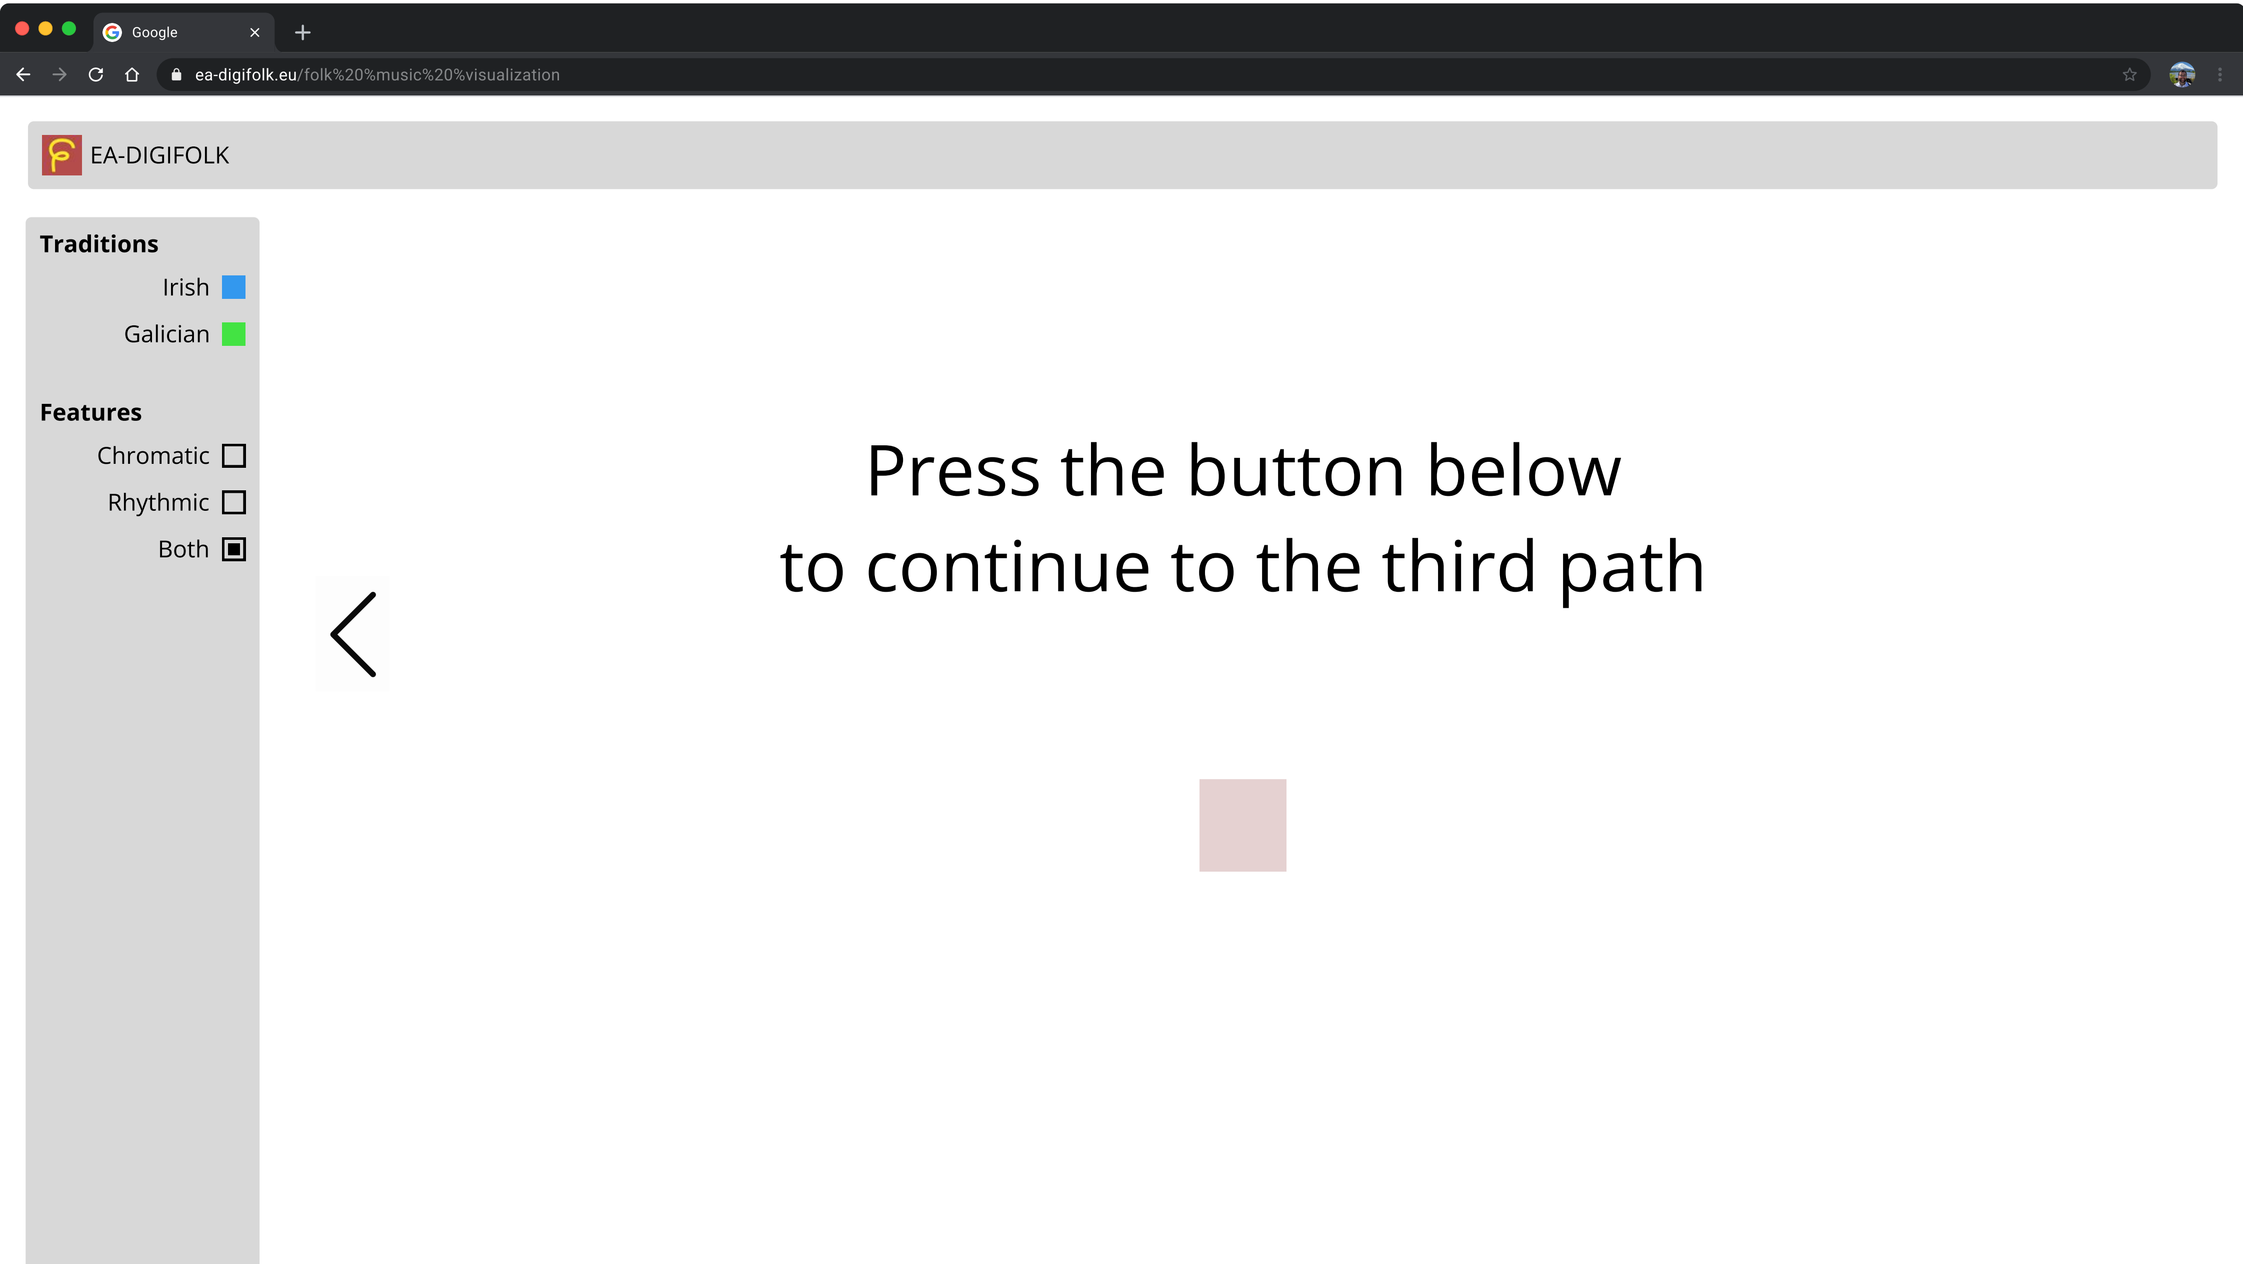
Google
ea-digifolk.eu/folk%20%music%20%visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
Press the button below
to continue to the third path

Google
ea-digifolk.eu/folk%20%music%20%visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
alalas
polka
slide
mazurcas
marchas
rumbas
polca
waltz
pasodobles
foliadas
jotas
valses
pasacorredoi
muineiras
jig
mazurka
march
hornpipe
barndance
reel
strathspey

Google
ea-digifolk.eu/folk%20music%20visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
muineiras
pasacorredoi
mazurka
jig
hornpipe
march
barndance
reel
strathspey
alalas
slide
polka
rumbas
polca
waltz
pasodobles
valses
foliada
jotas
marchas
mazurcas

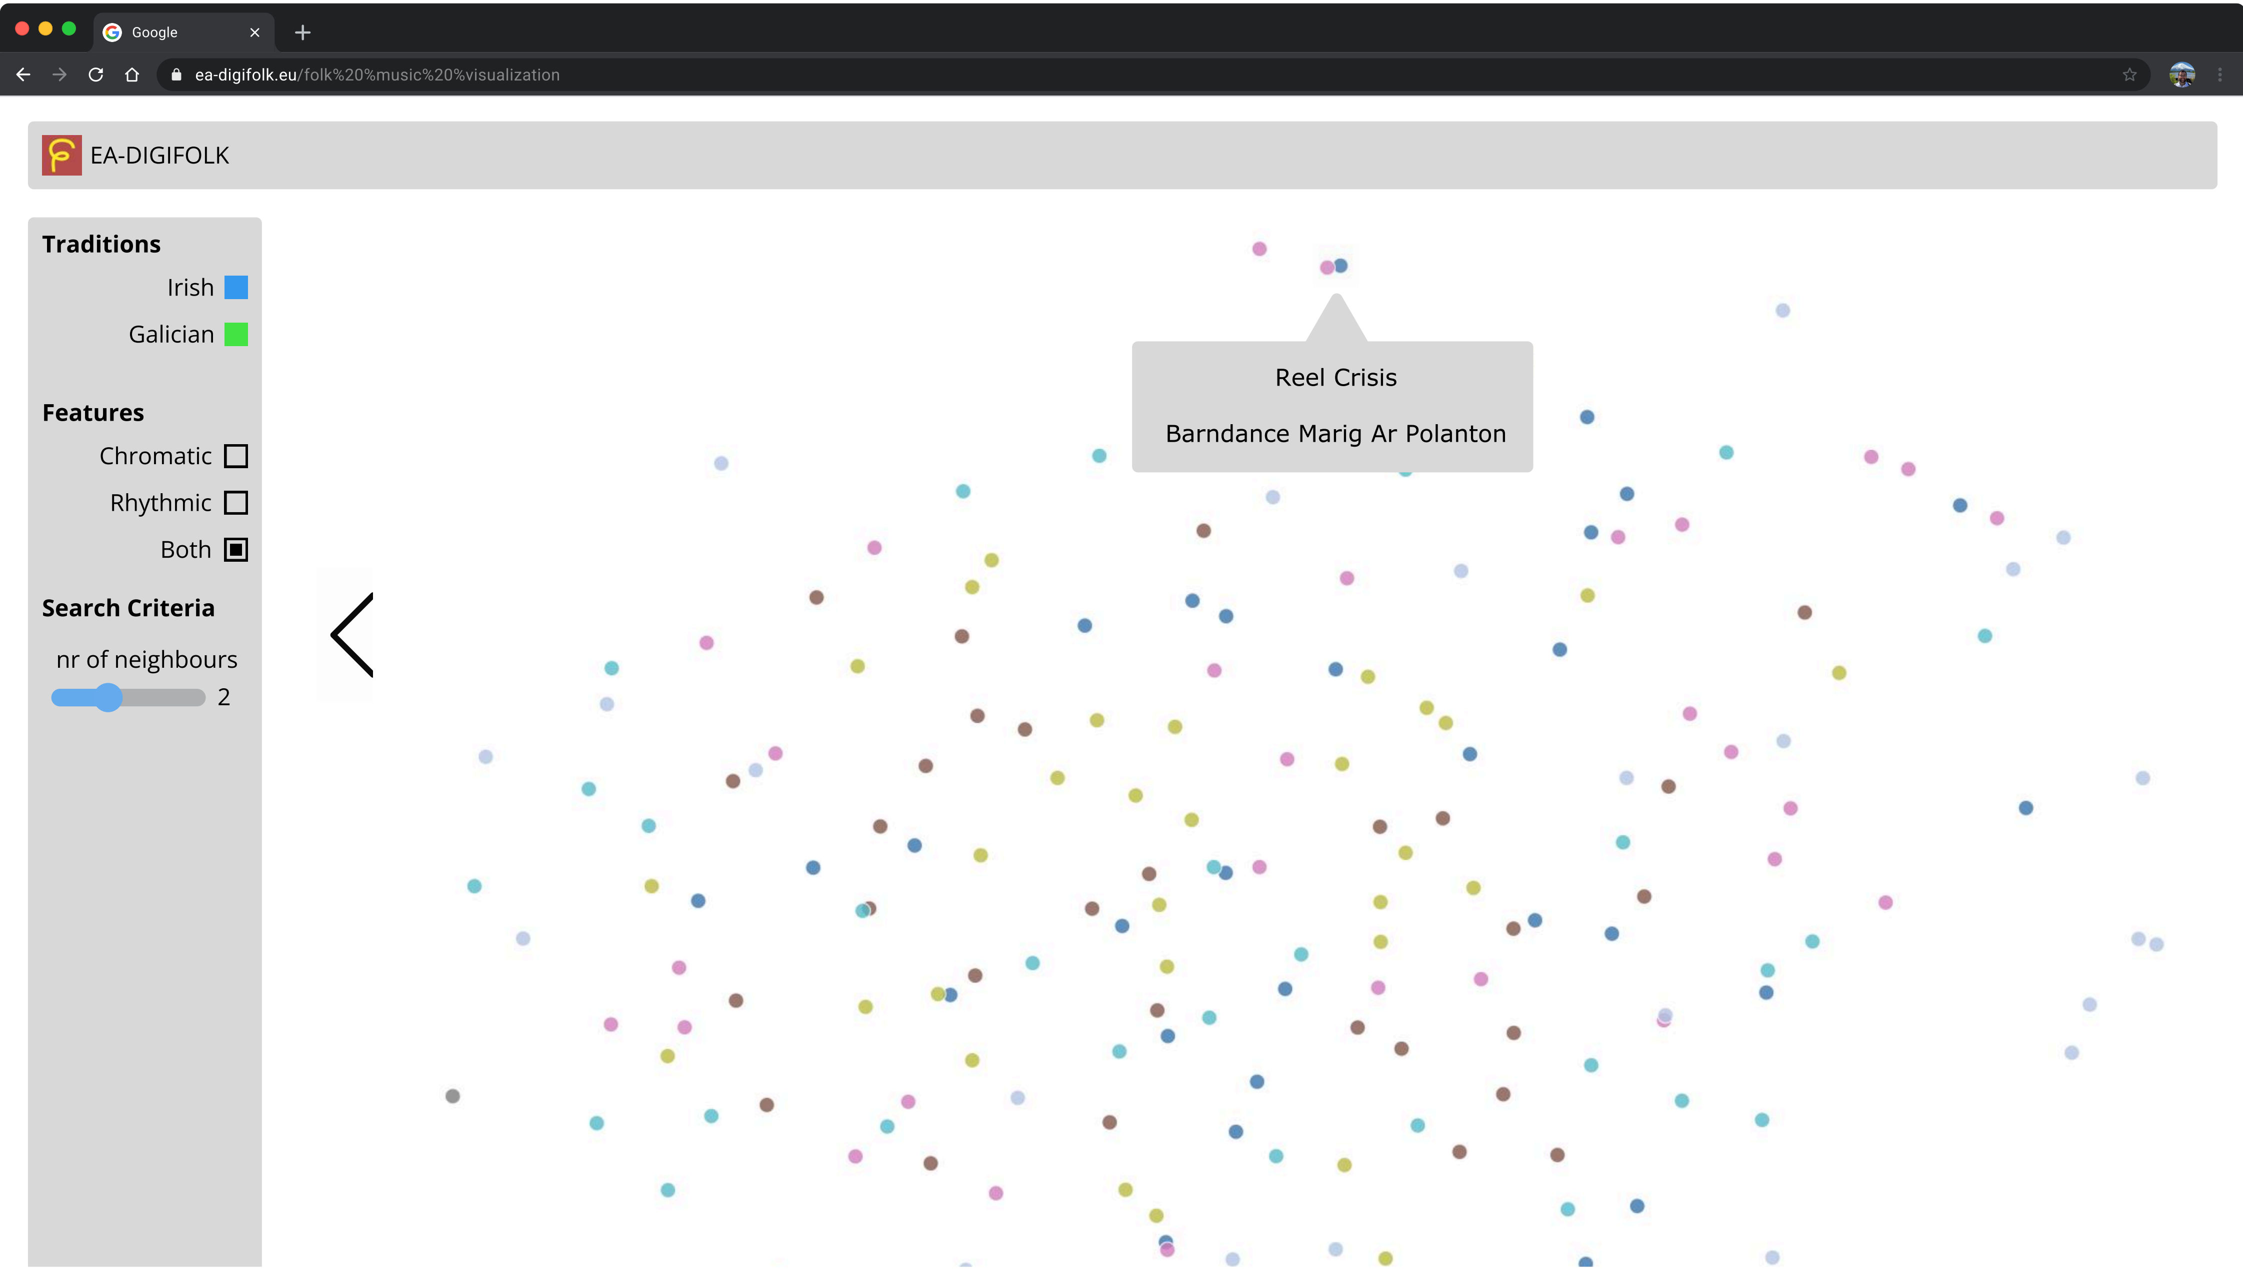
Google
ea-digifolk.eu/folk%20%music%20%visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
Search Criteria
nr of neighbours
2
Reel Crisis
Barndance Marig Ar Polanton

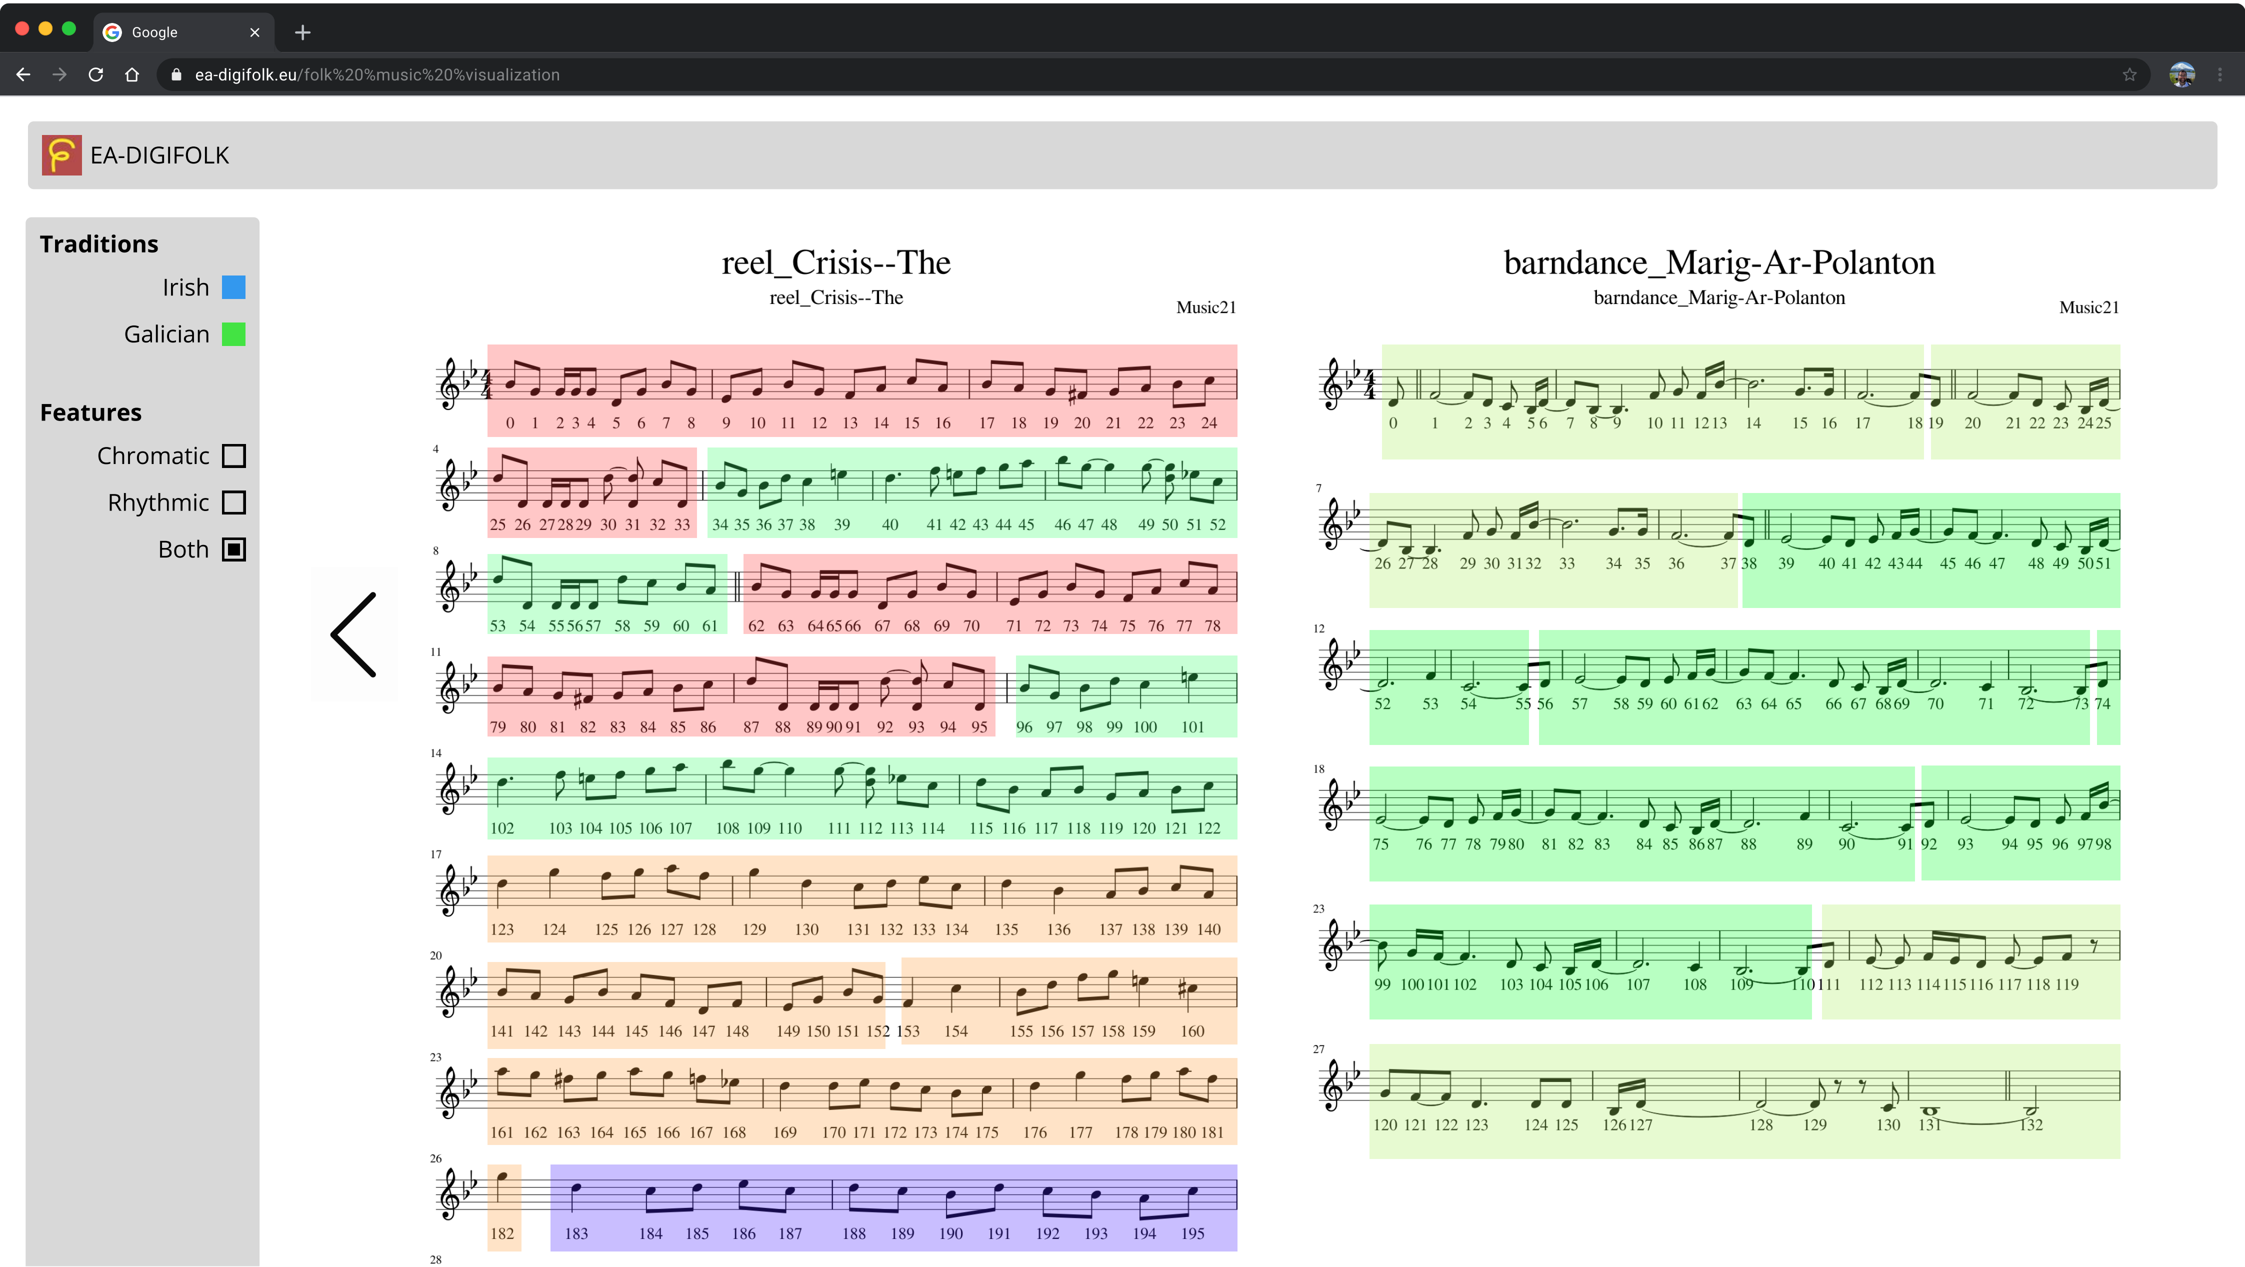

Google
ea-digifolk.eu/folk%20music%20visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
reel_Crisis--The
reel_Crisis--The
Music21
barndance_Marig-Ar-Polanton
barndance_Marig-Ar-Polanton
Music21

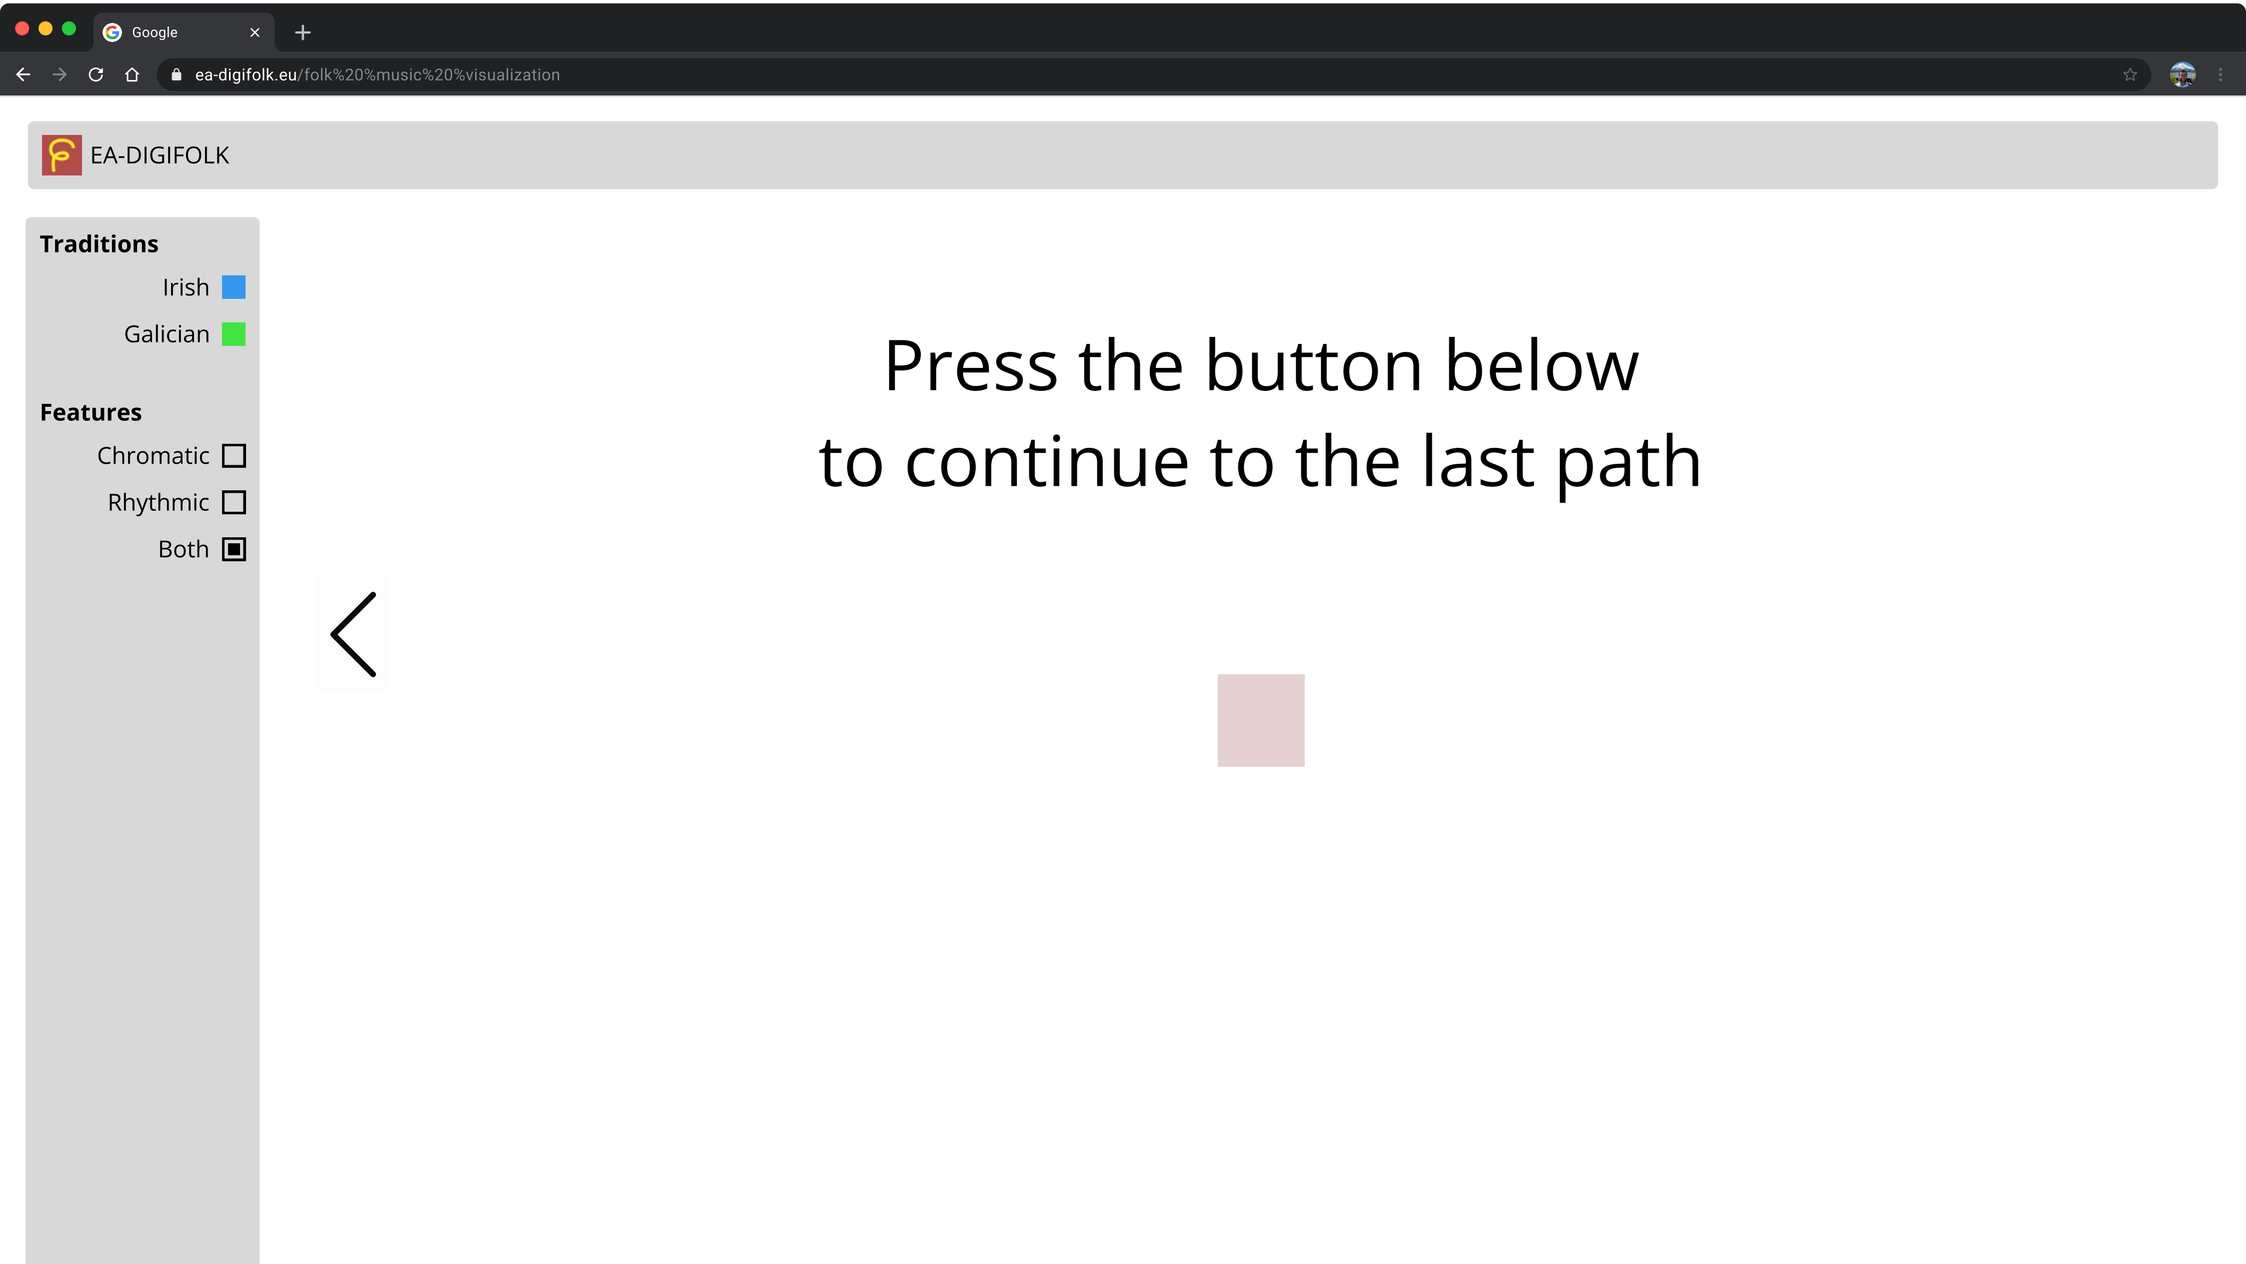
Google
ea-digifolk.eu/folk%20%music%20%visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
Press the button below
to continue to the last path


Google
ea-digifolk.eu/folk%20%music%20%visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
alalas
polka
slide
mazurcas
marchas
rumbas
polca
waltz
pasodobles
foliadas
jotas
valses
pasacorredoi
muineiras
jig
mazurka
march
hornpipe
barndance
reel
strathspey


Google
ea-digifolk.eu/folk%20music%20visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
muineiras
pasacorredoi
mazurka
jig
hornpipe
march
barndance
reel
strathspey
alalas
slide
polka
rumbas
polca
waltz
pasodobles
valses
foliada
jotas
marchas
mazurcas

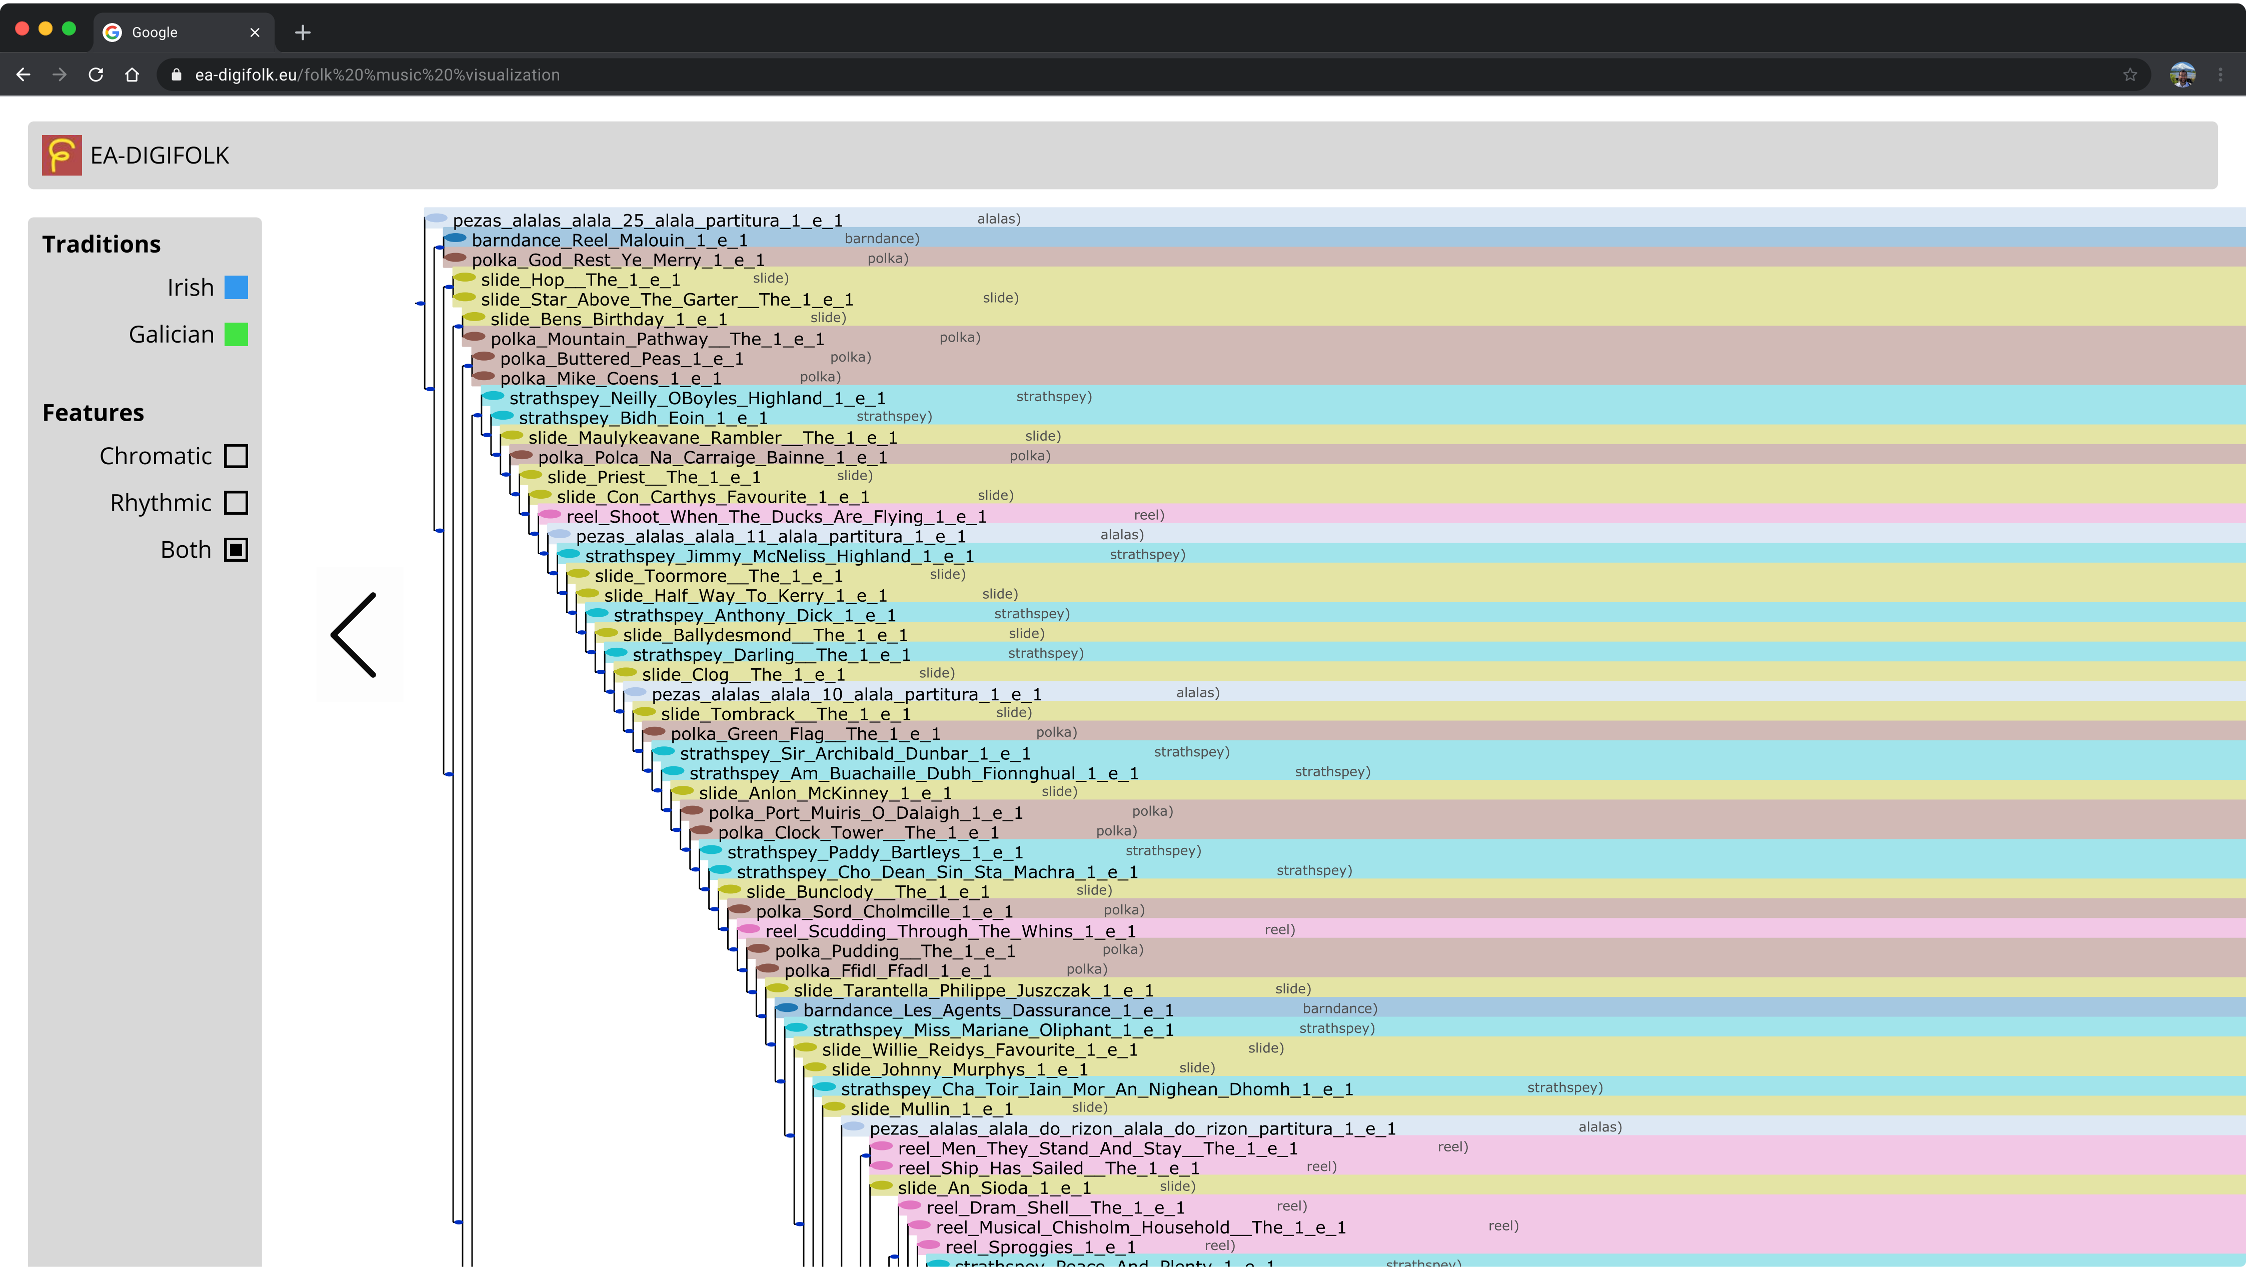
Google
ea-digifolk.eu/folk%20music%20visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
pezas_alalas_alala_25_alala_partitura_1_e_1 alalas)
barndance_Reel_Malouin_1_e_1 barndance)
polka_God_Rest_Ye_Merry_1_e_1 polka)
slide_Hop__The_1_e_1 slide)
slide_Star_Above_The_Garter__The_1_e_1 slide)
slide_Bens_Birthday_1_e_1 slide)
polka_Mountain_Pathway__The_1_e_1 polka)
polka_Buttered_Peas_1_e_1 polka)
polka_Mike_Coens_1_e_1 polka)
strathspey_Neilly_OBoyles_Highland_1_e_1 strathspey)
strathspey_Bidh_Eoin_1_e_1 strathspey)
slide_Maulykeavane_Rambler__The_1_e_1 slide)
polka_Polca_Na_Carraige_Bainne_1_e_1 polka)
slide_Priest__The_1_e_1 slide)
slide_Con_Carthys_Favourite_1_e_1 slide)
reel_Shoot_When_The_Ducks_Are_Flying_1_e_1 reel)
pezas_alalas_alala_11_alala_partitura_1_e_1 alalas)
strathspey_Jimmy_McNeliss_Highland_1_e_1 strathspey)
slide_Toormore__The_1_e_1 slide)
slide_Half_Way_To_Kerry_1_e_1 slide)
strathspey_Anthony_Dick_1_e_1 strathspey)
slide_Ballydesmond__The_1_e_1 slide)
strathspey_Darling__The_1_e_1 strathspey)
slide_Clog__The_1_e_1 slide)
pezas_alalas_alala_10_alala_partitura_1_e_1 alalas)
slide_Tombrack__The_1_e_1 slide)
polka_Green_Flag__The_1_e_1 polka)
strathspey_Sir_Archibald_Dunbar_1_e_1 strathspey)
strathspey_Am_Buachaille_Dubh_Fionnghual_1_e_1 strathspey)
slide_Anlon_McKinney_1_e_1 slide)
polka_Port_Muiris_O_Dalaigh_1_e_1 polka)
polka_Clock_Tower__The_1_e_1 polka)
strathspey_Paddy_Bartleys_1_e_1 strathspey)
strathspey_Cho_Dean_Sin_Sta_Machra_1_e_1 strathspey)
slide_Bunclody__The_1_e_1 slide)
polka_Sord_Cholmcille_1_e_1 polka)
reel_Scudding_Through_The_Whins_1_e_1 reel)
polka_Pudding__The_1_e_1 polka)
polka_Ffidl_Ffadl_1_e_1 polka)
slide_Tarantella_Philippe_Juszczak_1_e_1 slide)
barndance_Les_Agents_Dassurance_1_e_1 barndance)
strathspey_Miss_Mariane_Oliphant_1_e_1 strathspey)
slide_Willie_Reidys_Favourite_1_e_1 slide)
slide_Johnny_Murphys_1_e_1 slide)
strathspey_Cha_Toir_Iain_Mor_An_Nighean_Dhomh_1_e_1 strathspey)
slide_Mullin_1_e_1 slide)
pezas_alalas_alala_do_rizon_alala_do_rizon_partitura_1_e_1 alalas)
reel_Men_They_Stand_And_Stay__The_1_e_1 reel)
reel_Ship_Has_Sailed__The_1_e_1 reel)
slide_An_Sioda_1_e_1 slide)
reel_Dram_Shell__The_1_e_1 reel)
reel_Musical_Chisholm_Household__The_1_e_1 reel)
reel_Sproggies_1_e_1 reel)

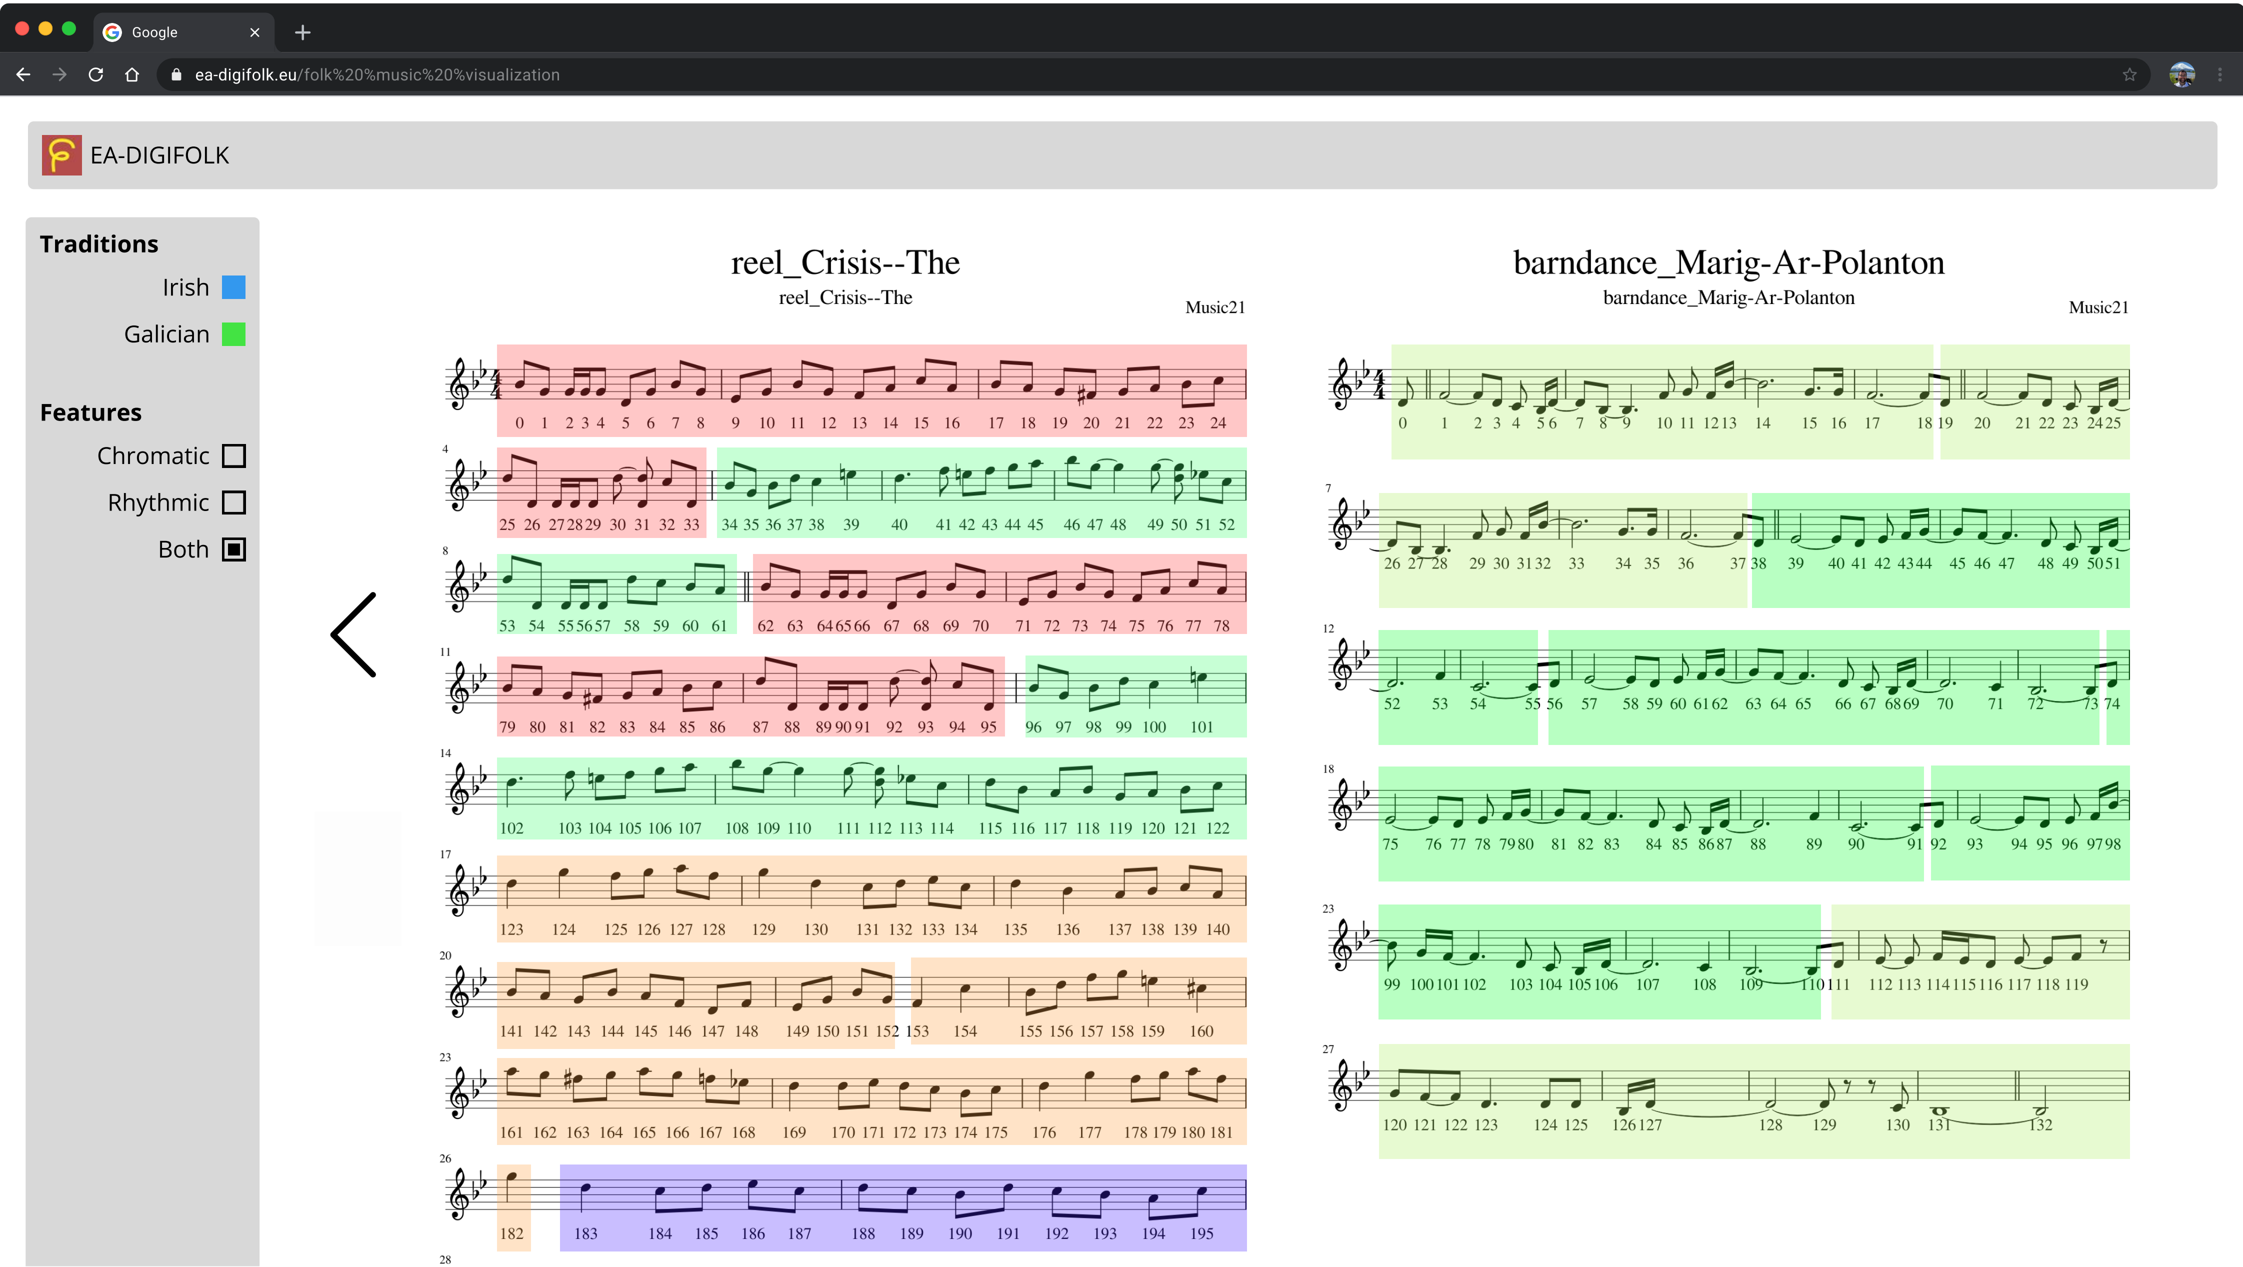
Google
ea-digifolk.eu/folk%20music%20visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
reel_Crisis--The
reel_Crisis--The
Music21
barndance_Marig-Ar-Polanton
barndance_Marig-Ar-Polanton
Music21

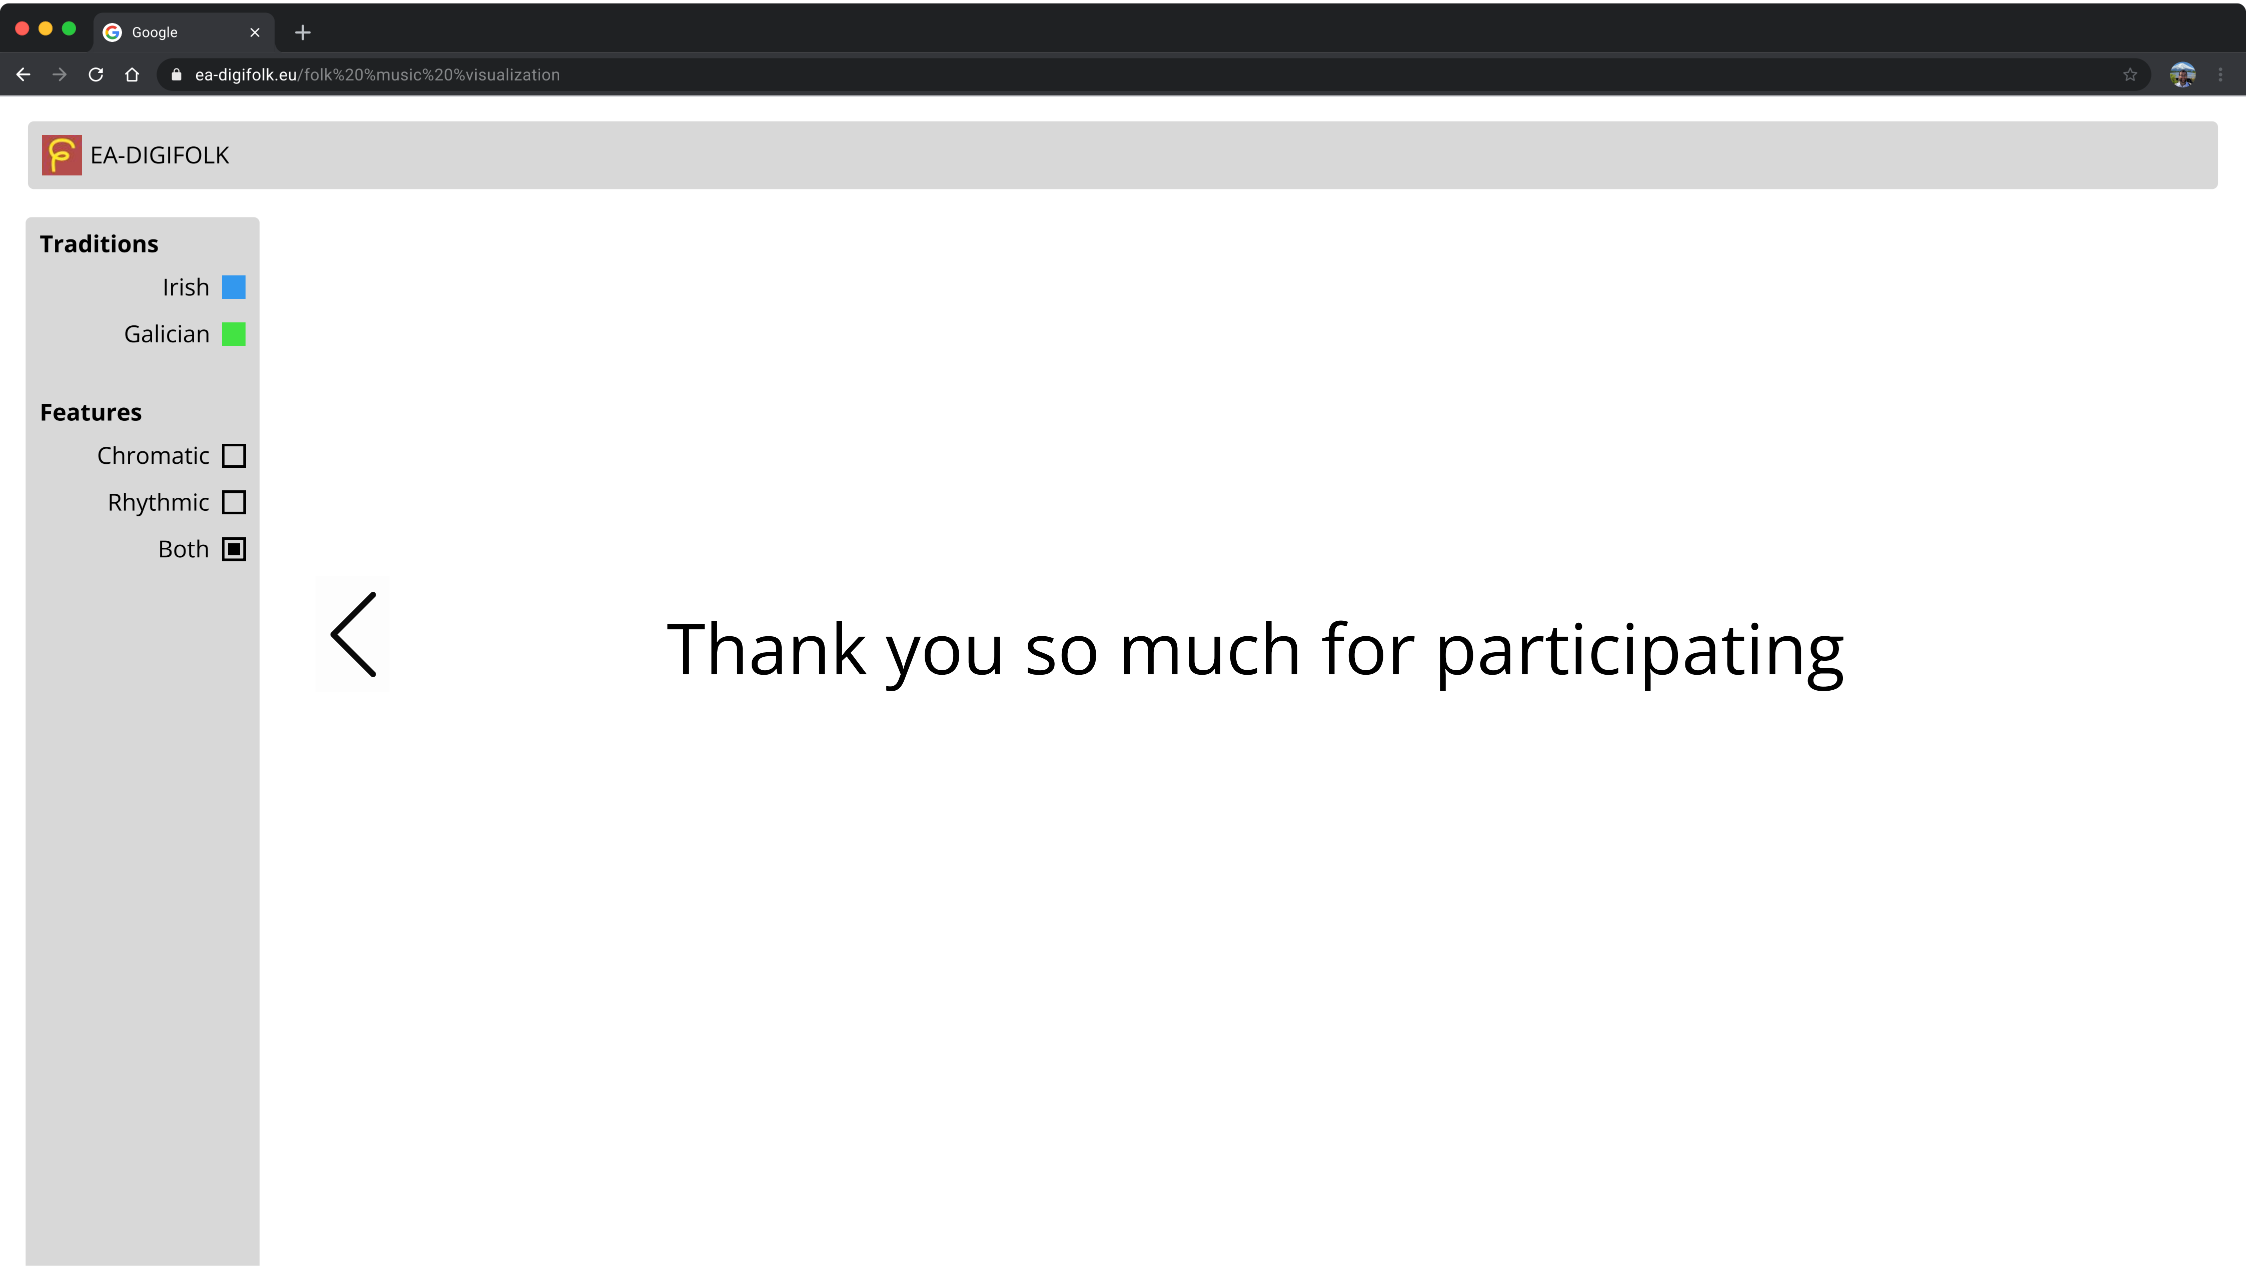
Google
ea-digifolk.eu/folk%20%music%20%visualization
EA-DIGIFOLK
Traditions
Irish
Galician
Features
Chromatic
Rhythmic
Both
Thank you so much for participating

Appendix C

Co-design Session Followup Form

C.1 Phylogenetic Approach to Folk Music Visualization

Thank you for participating in this evaluation of my folk music visualization tool. This web application is part of my Master's thesis project at the University of Porto, focused on visualizing structural relationships in Irish and Galician folk music traditions. The tool offers three main components:

1. Phylogenetic trees for exploring structural relationships between musical genres and individual pieces;

2. Interactive navigation between different tradition levels (from broad traditions to specific scores);

3. Score visualization and comparison for analyzing musical similarities and differences.

This project is part of the European EA-DIGIFOLK initiative, which aims to enhance access to and understanding of folk music traditions through digital tools.

What are Phylogenetic Trees?

In this application, phylogenetic trees are visual representations that show relationships between musical pieces based on their structural similarities:

-Branches represent evolutionary relationships between pieces or genres;

-Nodes (connection points) represent either individual scores or groups of related music;

-Distance between elements indicates their musical similarity—closer items are more structurally similar;

Just as biologists use phylogenetic trees to show evolutionary relationships between species, this tool uses them to reveal musical connections—showing how melodies, rhythms, and structures relate across different traditions and genres.

Your Feedback

Your expertise is invaluable in evaluating this tool. The question on the following page will help me understand how well the visualization achieves its goals and how it might be improved to better serve researchers, musicians, and folk music enthusiasts.

I would ask you to explore the tool for 1-2 minutes, and see some scores from a specific genre, and compare up to 3 scores from it.

Link to access the website: EA-DIGIFOLK Visualization

C.2 System Usability Scale

1. I think that I would like to use this system frequently.
   - 1
   - 2
   - 3
   - 4
   - 5

2. I found the system unnecessarily complex.
   - 1
   - 2
   - 3
   - 4
   - 5

3. I thought the system was easy to use.
   - 1
   - 2
   - 3
   - 4
   - 5

4. I think that I would need the support of a technical person to be able to use this system.
   - 1
   - 2

- 3
- 4
- 5

5. I found the various functions in this system were well integrated.
    - 1
    - 2
    - 3
    - 4
    - 5

6. I thought there was too much inconsistency in this system.
    - 1
    - 2
    - 3
    - 4
    - 5

7. I would imagine that most people would learn to use this system very quickly.
    - 1
    - 2
    - 3
    - 4
    - 5

8. I found the system very cumbersome to use.
    - 1
    - 2
    - 3
    - 4
    - 5

9. I felt very confident using the system.
    - 1
    - 2
    - 3
    - 4
    - 5

10. I needed to learn a lot of things before I could get going with this system.

- 1
- 2
- 3
- 4
- 5

C.3 Qualitative Evaluation

1. Did you participate in the previous co-design session?

1.2. How does the implemented visualization compare to the mockups we discussed? What works well and what doesn't?

1.3. Did you encounter any unexpected difficulties while navigating between traditions, genres, and scores?

2. Do you see practical applications for this tool in your work or research? If so, what specific scenarios?

3. Is the score visualization clear and functional for your needs? How might it be improved?

4. Any additional suggestions or critical feedback that would make this tool more valuable?

Thank you so much for participating!

Appendix D

Co-design Session Followup Results

System Usability Scale (SUS) Results

Participant	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10
1	5	1	4	1	3	1	4	1	4	1
2	5	4	2	3	4	4	1	3	4	5
3	1	2	5	1	1	4	1	5	1	1
4	4	2	4	3	4	1	4	3	3	2
5	4	1	4	2	3	2	4	1	3	1
6	2	2	3	1	3	1	2	3	3	2
7	2	2	4	3	4	1	3	3	3	2
8	3	2	3	4	4	3	5	3	3	4
9	3	2	4	1	3	2	4	2	3	2
10	4	1	5	2	5	1	5	2	4	5
11	4	2	5	1	5	2	5	1	5	1
12	2	1	5	1	3	1	4	1	2	3
13	3	1	5	1	5	1	3	1	3	1
14	4	1	5	1	4	2	4	1	4	4
15	3	3	4	2	3	3	4	2	4	2

Legend for SUS Questions:

1. I think that I would like to use this system frequently.
2. I found the system unnecessarily complex.
3. I thought the system was easy to use.
4. I think that I would need the support of a technical person to be able to use this system.
5. I found the various functions in this system were well integrated.
6. I thought there was too much inconsistency in this system.
7. I would imagine that most people would learn to use this system very quickly.
8. I found the system very cumbersome to use.
9. I felt very confident using the system.

10. I needed to learn a lot of things before I could get going with this system.

Qualitative Section

1. Did you participate in the previous co-design session?

- Yes: Participants 1, 11, 15;
- No: Participants 2, 3, 4, 5, 6, 7, 8, 9, 10, 12, 13, 14.

2. How does the implemented visualization compare to the mockups we discussed? What works well and what doesn't?

- **Participant 1**: The current version aligns quite well with what we discussed. The only improvements I would suggest are minor web design adjustments to make better use of space and enhance the overall layout. I'd also be interested in seeing some of the additional features we discussed gradually incorporated. Overall, it's a solid starting point.
- **Participant 11**: The visualization of the scores works well and it is interesting to have the possibility to download the MusicXML file, but it would also be interesting to be able to offer the scores encoded in the MEI standard and to offer the audio. On the other hand, although elements such as the shape of the song are available in the filters, many of them are not available in the display.
- **Participant 15**: The layout is nicer and the flow of the hierarchy is easier to follow.

3. Did you encounter any unexpected difficulties while navigating between traditions, genres, and scores?

- **Participant 1**: No
- **Participant 11**: It is not clear to me the degree of similarity between scores. It might be interesting to specify what makes the nodes related.
- **Participant 15**: There were quite a few 'score not found' errors.

4. Do you see practical applications for this tool in your work or research? If so, what specific scenarios?

- **Participant 1**: Yes, I see practical applications for this tool in my research. It would be especially useful for analyzing large collections of scores without having to inspect each one manually. It could also help identify examples and cases of scores that are similar to the ones I'm working with, which is valuable for comparative analysis and pattern discovery. Additionally, it could support more efficient dataset exploration and hypothesis generation in symbolic music research. As a musician, it would also make it easier to find pieces similar to the ones I'm currently playing, which is helpful for repertoire expansion and stylistic exploration.
- **Participant 2**: Yes. I have searched and written about generative analysis in another kind of music in my dissertation.
- **Participant 3**: No

- **Participant 4**: Yes, it could support educational initiatives by offering interactive visualizations that make it easier to understand the historical and structural relationships between different folk traditions.

- **Participant 5**: It seems like an intuitive way to look at similarities between the different styles. It's interesting to see the existing relationships between styles and it makes me look at them in a way I maybe haven't before.

- **Participant 6**: (a bit of background, I'm an Irish folk musician, and I do research in protein bioinformatics and music evolution) I'm working on something slightly similar, so I have different ideas about what is useful. I work directly with data, so I do computational analyses. The idea of visualization for score comparison is cool, and definitely something I'd be interested in. I would be more interested in a feature where I can search for, or input tunes, to look for similar ones, and compare how/why they are similar. I can't think of what I would use the phylogenetic tree for, at least in the way it's shown. My first thought was to look at the end-nodes, to compare similarity, but I didn't notice the tunes being that similar (compared to what I expect, given my knowledge of Irish folk music).

- **Participant 7**: Historic cultural and music theory focused scenarios.

- **Participant 8**: Yes. it can gives a better understanding of each tradition with all complex variations.

- **Participant 9**: I see that the model is useful for my work. In my work as a music teacher it is very important to select repertoire based on musical and pedagogical criteria. It is interesting to find similar songs with which to work with the students and, in this way, also to contextualize the musical practice socially.

- **Participant 10**: I use a similar tool for teaching electronic music, Ishkur's Guide to Electronic Music. It helps my students know what musics are out there and how they're related to each other both musically and in history.

- **Participant 11**: Yes, to try to find similarity between musical pieces quickly (as I am research in the EA-Digifolk project, it would make a lot of sense to use this tool).

- **Participant 12**: Yes, with a larger dataset I can see it working for comparing tunes.

5. Is the score visualization clear and functional for your needs? How might it be improved?

- **Participant 1**: Yes, the score visualization is clear and functional for my needs. It would be useful to have the option to select and compare scores from different genres or traditions. A search bar would also help when looking for a specific piece, instead of browsing through the tree. Additionally, showing relationships between scores from different genres or traditions could make the tool more insightful.

- **Participant 2**: No. More examples and more cases needed.

- **Participant 3**: No. It is completely unclear how this tool works. Clicking different options appears to do nothing.

- **Participant 4**: The score visualization is generally clear and functional, especially for identifying structural similarities and differences between musical pieces.

- **Participant 5**: Seems fine.

- **Participant 6**: The comparison is missing something fundamental, which should be in principle solvable, but it's not clear how easy it would be. The most important addition that I think will eventually be added, is something like what you have for sequence alignment of biological sequences. It's not clear how to do this on a musical score, but there should be some way of highlighting which parts are same and different – this means annotating which parts are aligned (if local alignment), where there are gaps, and where there are mismatches (and how big the mismatch is).

- **Participant 7**: Sound and with theoretical assistance.

- **Participant 8**: Yeah however if we can play the song and hear the differences or some further explanation focusing on main characteristics and origin and similar patterns and families would be more intuitive in understanding the bigger picture.

- **Participant 9**: I have already commented above, it would be interesting to provide the MEI files of the scores and also the context of the songs. Traditional music is usually collected with annotations about what time of the year it is sung, where...

- **Participant 10**: Some of the scores have a single measure taking up an entire staff (jig: Catch Me If You Can), while other staves have multiple measures. This inconsistency in how many measures per line would be confusing to play from. Some of the scores give me the error: Error Loading Score... Cannot read properties of undefined (reading 'substring') Please make sure your MusicXML files are correctly placed in one of these directories.

- **Participant 11**: It would be interesting to make the level attributes stand out on the score (e.g., where phrases are similar). It would also be interesting to be possible to hear the scores.

- **Participant 12**: It works very well with a small dataset, but I would be concerned that with larger numbers it might be harder to follow.

6. Any additional suggestions or critical feedback that would make this tool more valuable?

- **Participant 1**: Although this is primarily a web-based application, it would be helpful to ensure that it is also optimized for mobile devices. Improving mobile responsiveness would make the tool more accessible and convenient in different usage contexts.

- **Participant 4**: One suggestion to increase the tool's value would be to allow users to upload and analyze their own datasets, enabling broader use across diverse music traditions.

- **Participant 5**: When we click a style and see all the available songs it gets a bit cluttered and confusing. Some text overlaps with the tree lines and with other text, and some song titles have this "1 E 1" at the end, idk what that means. I think it would also be interesting to see why things are related the way they are as I explore the trees, but I have no concrete ideas of how that might look like.

- **Participant 6**: The most important addition that I think will eventually be added, is something like what you have for sequence alignment of biological sequences. It's not clear how to do this on a musical score, but there should be some way of highlighting which parts are same and different – this means annotating which parts are aligned (if local alignment), where there are gaps, and where there are mismatches (and how big the mismatch is).

- **Participant 8**: 1. Generally maybe it would be good to present the short form of each song to make the graph more minimalistic and not too crowded. 2. If I click on the nodes with no label, it gives an error. It would be better to fix it in case of accidental click. 3. If we can have very short info for each song or family on the tree level, it would be very enjoying to navigate before opening the score. Maybe implementing it on a map and highlight that region might be even more eye-catching.

- **Participant 10**: This seems like it has the potential to be useful, but there aren't clear explanations of the Features and Level controls. I also don't know what each pre-score node means. How are the songs on the same branch related to each other in the various Level views? It would be good to have an explanation associated with each node in the tree, not just the terminal nodes.

- **Participant 11**: There was a bug for me. When I clicked on the Valse node, a black screen was showing but no score (Ubuntu, Mozilla browser).

- **Participant 12**: The ability to compare across rhythms and across traditions would be very useful.