JazzDAP: Collaborative Research Tools for Digital Jazz Archives

Kevin Allain
Kevin.Allain.2@city.ac.uk
City, University of London
London, United Kingdom

Tillman Weyde
T.E.Weyde@city.ac.uk
City, University of London
London, United Kingdom

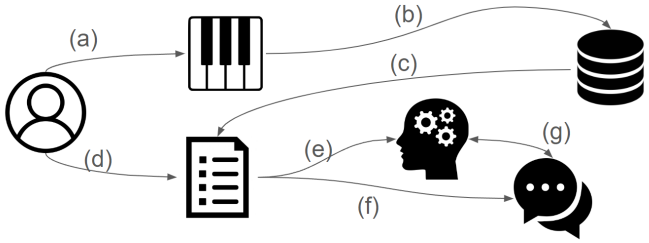

Figure 1: Flow of JazzDAP system use. The user (a) enters a melodic pattern and filters that are (b) passed to a database. The results are (c) displayed to the user. The user explores them (d) to generate insights (e) that can be saved into workflows, i.e., aggregates of objects of interest and user notes. Discussions are enabled for elements of interests of these results (f) as well as for workflows (g) to enable collaboration between users.

ABSTRACT

This paper introduces a novel web platform designed for exploration, analysis, and collaboration in the jazz music domain called JazzDAP. Our platform integrates advanced music information retrieval techniques with user-friendly interfaces, tailored for musicologists, archivists, and jazz enthusiasts. The platform employs a contour based algorithm for pattern recognition, enabling users to search for specific musical motifs, with filters based on metadata, e.g. artist, location and year of recording. Users can listen to audio sections or MIDI excerpts from the matches and delve into detailed metadata, including the years of recordings, prevalence of specific patterns, and information about the artists associated with them. Visualizations aid in uncovering trends, evolution, and connections in the development of jazz. A key innovation of our platform is the introduction of *workflow* objects, allowing users to save elements of interest accompanied by notes, named workflows, and engage in collaborative discussions. Users can use workflows to annotate, share insights, and communicate with each other, fostering a community-driven exploration of jazz music. This collaborative aspect enhances the platform's utility for researchers and enthusiasts alike, aiming to create a dynamic environment for the exchange of knowledge and discoveries. This paper outlines the platform's structure, highlights its key features, and presents preliminary user feedback. We believe that our work opens new avenues for the exploration and understanding of jazz music, offering a valuable resource for researchers, archivists, and enthusiasts passionate about the intricate patterns that shape the genre.

DLfM 2024, June 27, 2024, Stellenbosch, South Africa

ACM ISBN 979-8-4007-1720-8/24/06
https://doi.org/10.1145/3660570.3660572

KEYWORDS

Music Information Retrieval, Jazz Music Archives, Digital Musicology, Pattern Recognition in Music, Music Metadata Exploration

ACM Reference Format:
Kevin Allain and Tillman Weyde. 2024. JazzDAP: Collaborative Research Tools for Digital Jazz Archives. In *11th International Conference on Digital Libraries for Musicology (DLfM 2024), June 27, 2024, Stellenbosch, South Africa.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3660570.3660572

1 INTRODUCTION

In the domain of jazz music research, the promise of a "big data" and "artificial intelligence" offers a new approach to understanding the intricate fabric of this genre that evolved over the last one-hundred years. Projects that ran in the last decade, like Digital Music Lab [20], J-DISC [4], Jazzomat [13], and Dig That Lick [5], which focused on infrastructure, metadata and tools to analyse digital music collections and specifically jazz recordings, beginning to address regarding the prevalence and evolution of melodic patterns and licks over time. How much of jazz improvisation consists of stock patterns played verbatim, and how do these patterns vary across different stylistic contexts and age groups of musicians? These inquiries underscore the potential of a data-driven approach in shedding light on the nuanced dynamics within jazz performance traditions. While existing research provides valuable insights, our endeavour goes beyond data analysis. We integrate the element of collaboration, recognizing the importance of collective engagement and discourse in enriching our understanding of jazz. Our research and development is driven by several key challenges in jazz music scholarship. Accessibility of the archival materials online remains a significant hurdle, hindering researchers' ability to explore the rich historical tapestry of jazz. Moreover, navigating the complex web of musical context presents a formidable task, compounded by the inherent difficulty in searching for specific musical elements within vast repositories of audio recordings. We specifically focus

on melodic patterns, recognizing both their significance in jazz improvisation and the challenges associated with their identification and analysis. Furthermore, the absence of robust solutions for facilitating communication and collaboration among musicologists exacerbates these challenges, highlighting the need for innovative platforms that foster dialogue and knowledge exchange. In response to these challenges, our approach combines MIR technology with a collaborative framework to develop a novel platform for exploring jazz music archives. Through the integration of advanced search capabilities, interactive visualization tools, and collaborative features, our platform seeks to democratize access to jazz archival materials and facilitate meaningful discourse among musicologists and enthusiasts.

2 RELATED WORK

The focus of our project is on digital analysis of jazz libraries, and we discuss here prior related work.

The J-DISC project[1][4] was focused on digital jazz discography. It provided a dataset detailed, searchable metadata encompassing 2,711 unique sessions involving 3,744 artists. Analyses reveal trends in location, time period, and recording studios utilized. J-DISC provided metadata about jazz recording sessions, as they focus on collaboration patterns and their relation to studio and location.

In the Jazzomat project[2] [13] a comprehensive suite of tools and datasets was developed for understanding the structure and characteristics of jazz. The Weimar Jazz Database [12], provides researchers with a rich repository of high-quality transcriptions and annotations of improvised jazz solos. This allows for identification of musical patterns and improvisational techniques employed by musicians. While Jazzomat primarily focuses on computational analysis, the techniques and datasets provide the basis for projects like ours that seek to enhance user interaction and collaboration in exploring musical patterns.

The Dig that Lick project[3],[5] was an effort to automate aspects of the analysis of jazz recordings. Analysis of individual licks and large-scale pattern exploration, automatic pattern identification, comparison were explored. The tools developed for melodic pattern search, including algorithms applicable to different musical styles and strategies for tracking pattern evolution between musicians over time. Dig That Lick primarily focused on isolated analysis, lacking built-in collaborative tools for interactive exploration and knowledge sharing. This gap motivated, in the development of this project, the aim to support collaborative jazz pattern analysis with intuitive search, visualization, and annotation functionalities.

A common thread among these projects are their efforts to provide visualizations of the data, as musical data records are complex and rich. In order to analyse music in digital datasets, digital tools are needed to realise the idea of *distant reading* introduced by Moretti [10, 11]. While a number of methods and tools have been developed for text, [see 6], less has been done for music in this respect. The Digital Music Lab project applied big data visualisation techniques to music chord progressions [7]. The Jazzomat visualizations help understand the relationships between patterns in solos, but do not provide user accessible system. Dig That Lick included visualizations to display connections between groups of metadata out of the output of melodic patterns searches [3]. Another interesting example is the usage of Tableau by Matuszewska [9]. While limited to a single collection (Oskar Kolbeg's Opera Omnia), Matuszewska's visualization provides a rich array of graphs and dynamic filtering of data selection that enable musicologists to explore different aspects of the music collection.

A novel aspect of this work is the support for collaboration. The value and challenges of collaboration in digital humanities have been discussed, e.g. in [18] and [2], but there are not many platforms designed for digital humanities, especially not for digital musicology. Well-known collaborative platforms exist in other disciplines, specifically in computing GitHub [19] and Jupyter notebooks [14] are well-known examples, with millions of projects that have benefited from collaboration. We have taken inspiration form these platforms for the design of JazzDAP. GitHub is a software developer platform that allows for collaboration, community building, knowledge sharing and increases efficiency by keeping projects online, and enabling to make them open source. Jupyter allows for development of a narrative linked to a coding project, enabling interdisciplinary insights and increased accessibility and usability. Jupyter notebooks can include visualizations, increasing the value of their narrative. However, neither GitHub nor Jupyter are suitable to enable audio analysis or research into melodic patterns for non-programmers. The aim of JazzDAP is to build on the foundation laid by previous research such as Dig that Lick and expands its scope by developing tools that foster shared exploration and collective understanding of melodic patterns within the jazz community.

3 THE JAZZDAP PLATFORM

3.1 Data and technical structure of the platform

The JazzDAP platform is web-based and provides access to music recordings and their metadata through search, filtering and sharing *workflows*. Currently, we hold musical records from two sources: from the Rutgers University – Institute of Jazz Studies (the Benny Carter and Benny Goodman collections), and from the Scottish Jazz Archive, a regional data collection, their complete holdings. Transcriptions of these records are generated using specialised deep learning and signal processing methods [1, 16].

The technical architecture consists broadly of three parts:

- Front-end platform for user interaction: our website front-end is developed with JavaScript and React.js. MIDI sound is generated using the Tone.js library. Connection to the back-end is created with Axios.js.
- Back-end: We developed the back-end as a REST API with Node.js. The back-end also incorporates MP3 files with the music recordings. When a user wishes to listen to an audio section, these files are either dynamically split with a Python script or loaded from previously executed splits.
- The database where all data of the collections is stored is build with MongoDB and data is curated using Python.

3.2 The search interface

The JazzDAP platform is designed with two parts: the search interface that is used to find information based on musical patterns and

[1]The project homepage http://jdisc.columbia.edu is no longer reachable.
[2]https://jazzomat.hfm-weimar.de.
[3]http://dig-that-lick.eecs.qmul.ac.uk, see also [17] for a review of the project and its outputs.

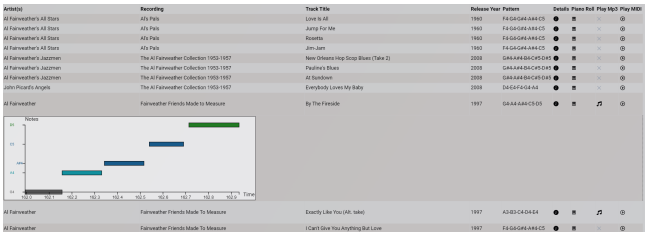

Figure 2: The user interface displays output of a melodic pattern search. A piano roll is displayed for a match in a track called 'By The Fireside'.

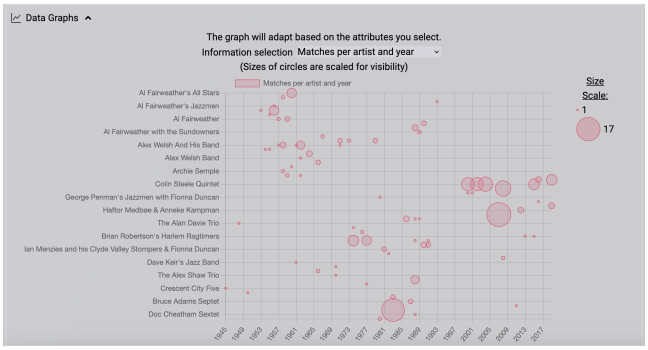

Figure 3: This bubble chart displays matches according to artists over the years. The size of the bubbles indicates the number of matches. This graph allows considering contextual evolution of pattern usage for an artist.

their context, and the workflow interface, designed to aggregate information of interest and enable collaboration by sharing, searching and discussion workflows. In this section, we detail the search interface.

3.2.1 Search input. Users can search through the database by entering a series of notes, or patterns. Notes are directly identified from MIDI that the previously mentioned algorithms transcribed from the audio recordings. Alternatively, a pattern is defined by a series of intervals between the notes. A search is made in two steps. First, series of notes are selected that share the same contour as the input (fuzzy intervals between subsequent pitches categorized in 9 groups). This step is quick as fuzzy interval scores are pre-computed for all note sequences in the database. Second, from this selection, the similarity of note sequences is evaluated based on a modified Levenshtein distance [8], employing each interval between two notes like a letter. For each note sequence that is a potential match, we only return those that have a distance to the input notes lower than a set threshold, which can be modified by the user.

The JazzDAP search provides filters in the melodic pattern search. These filters can be separated into two categories: melodic patterns filters and contextual filters. The melodic patterns are used to define the difference that is acceptable for a series of notes found in the database from a search, while the contextual filters are used prior to a search to reduce the number of potential matches. The contextual parameters are based on the following attributes: Track title, Name of the artists, Name of the recording (can be a published album, but in the currently used collections the recordings were never released publicly), Location of recording (the textual description of locations as provided, there is no standardised encoding of locations, yet, in the system), Year of recording, and Producer.

3.2.2 Search outputs. Following a search entered by a user, we display the note sequences that matched the search pattern as a table. For each melodic pattern match, further information about the record can be displayed once selected. For each entry in the match table, the matched note sequence can be presented with a piano roll by clicking a button. Other buttons allow playing the matching audio section and playing a MIDI rendering of the matching note sequence found in the database. The structure of the display is shown in Figure 2.

To help the user explore the context of the melodic pattern matches in the database, we provide a facility for creating graphs that display numerous aspects of the metadata. The platform allows displaying the matches over the following dimensions, using bar charts or bubble charts:

- Matches per year
- Matches per recording
- Matches per track
- Matches per artist
- Matches per pattern
- Matches per artist and year
- Matches per pattern and year

The match count per pattern is useful when the user sets a boundary on pattern to search. A bubble chart of the the matches per artist and year is shown in Figure 3).

3.3 Workflows and collaboration

3.3.1 Collaborative music search and analysis. Previous efforts to provide access to musical archives aimed at increasing the analytic capabilities provided to the user, but did not offer collaborative tools. The JazzDAP project was developed with this idea from its inception. We are not aware of prior work in this direction, and therefore we started with a broad approach, and set our platform to allow for the development of a narrative like Jupyter notebooks, and with tools to discuss and comment production for other users, like GitHub. The aim was to make the platform flexible enough to let users define themselves the objectives behind their usage of the platform, and work collaboratively without restrictions. Collaboration in the JazzDAP platform is based on the notion of a *workflow* object, where users can store items of interest found during their search and provide a narrative with their findings, and annotations and comments, such as open questions to the community. A workflow can be made public, so that it is visible to other who can also leave comments, enabling a discourse between users.

3.3.2 Annotations and Comments. Users can annotate their workflows and searches and objects found from the database, once items of interests are saved in a workflow. The purpose of this approach is to encourage discussions between users: once a user searches

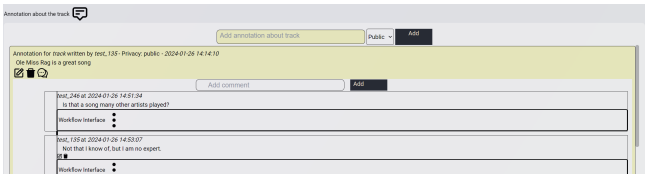

Figure 4: An example of a public discussion between two users. The first user makes an annotation about a track, and made it public to enable discussions. Following the creation of the public annotation, another user writes a comment related to the annotation, and the first user engages in the discussion by leaving another comment.

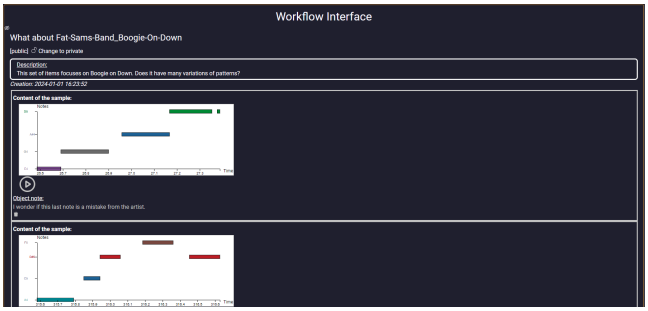

Figure 5: An example of a workflow, where the user saves several matching not sequences of interest and writes annotations to them. Here, the user focuses on the song 'Boogie On Down' and search patterns where results vary much. The matches that were found in the search interface are displayed in piano rolls and the MIDI can be played back.

through workflows, we expect that they have a specific rationale behind their approach, while melodic pattern search is more likely to result in diverse findings. Users can save annotations or comments in workflows, presented in the next subsection. Figure 4 illustrates a set of comments underneath an annotation.

3.3.3 Workflows. Workflows store and aggregate all objects of interest to users. The design of the workflows is inspired by Jupyter Notebook [14], where users can write details about code they have written in their project in a linear structure. A workflow has a title and and a description, which can be derived from an annotation or comment, to describe the intention behind the search, filtering and aggregation. of aggregating of objects of interest they encountered in the search interface, followed by the objects of interest they wish to save, as illustrated in figure 5. Each object of interest can be accompanied by some notes. According to users' preferences, the status of a workflow can be set to public or private.

The range of methods of analysis in musicology is wide [15] and thus the workflow was designed with the objective to save diverse objects of interests and write notes about them, leaving the users the freedom on how they intend to present their findings and reflections. Our hope is that this flexibility can allow musicologists and music enthusiasts to integrate the JazzDAP workflow into their investigative process.

4 USE CASE

To illustrate how JazzDAP can be used to collaboratively enhance users' searches, we discuss a hypothetical scenario where two users can find further elements of interest, by utilizing each other's respective findings. In this scenario, two users, Alice and Bob, with different searches, find they are looking at similar objects.

- Alice's search returns matches filtered to select a specific artist. She then looks for the matches of that artist over years. She saves elements of interest in a workflow designed to consider the music evolution of that artist over time.
- Bob searches for a different pattern, but focuses on a specific location. Bob then notices the same artist of interest, and search for workflows with this artist. Amongst the workflows returned, he reads Alice's.
- The workflow of Alice is interesting to Bob as he sees another recurring pattern for that artist who was prevalent in the location he was interested in. He sets an annotation to Alice's workflow, and communicates about the artist and the global context of the music in that era and geography with Alice.

This scenario is defined by two users who have set objectives in mind, but the approach also applies for more exploratory searches without a specific objective that can result in productive discussions, where insights or potentially more specific questions evolve.

5 FEEDBACK AND FUTURE WORK

We presented the JazzDAP platform at the Jazz Congress 2024 in New York City, as well as to a panel of musicologists and archivists during a workshop at Rutgers University, Newark, NJ. The presentation of the prototype allowed for focused discussion on what the features are that users would appreciate in an interface for a jazz archive. Listening to the audio of matched patterns and finding information about the artists emerged as features of interest.

6 CONCLUSION

Music analysis can benefit from access to digital jazz archives' collections, and digital tools are necessary to process the volume of information. With JazzDAP, we introduce a novel approach to enable collaborative analysis of digital jazz recordings using workflow objects. Discussions with field experts indicated the value of our approach and the promise of future extensions of the JazzDAP platform. Tasks for future work are

- Integration of data and metadata available on other platforms (e.g. Wikipedia)
- Enrichment of the search by including timing in input, and display of more contextual information (e.g. chord changes)
- Direct input of MP3 or MIDI to search the database
- Enabling users to save their own MP3 or MIDI, e.g. in a shared workflow.

We believe that collaboration will foster the discourse in jazz research and by extending the features and continuing to gather input from users and other stakeholders, the user interfaces can contribute further to the accessibility of digital jazz archives.

REFERENCES

[1] Dogac Basaran, Slim Essid, and Geoffroy Peeters. 2018. Main melody extraction with source-filter NMF and CRNN. In *19th International Society for Music Information Retreival*.

[2] Anela Chan, Richard Chenhall, Tamara Kohn, and Carolyn Stevens. 2017. Interdisciplinary collaboration and brokerage in the digital humanities. *Digital Humanities Quarterly* 11, 3 (2017).

[3] Klaus Frieler, Frank Höger, Martin Pfleiderer, and Simon Dixon. 2018. Two Web Applications for Exploring Melodic Patterns in Jazz Solos.. In *ISMIR*. 777–783.

[4] Yun Hao, Kahyun Choi, and J Stephen Downie. 2016. Exploring J-DISC: Some Preliminary Analyses. In *Proceedings of the 3rd International workshop on Digital Libraries for Musicology*. 41–44.

[5] Lucas Henry, Klaus Frieler, Gabriel Solis, Martin Pfleiderer, Simon Dixon, Frank Höger, Tillman Weyde, and Helene-Camille Crayencour. 2021. Dig that lick: Exploring patterns in jazz with computational methods. *Jazzforschung/Jazz Research* 50 (2021).

[6] Stefan Jänicke, Greta Franzini, Muhammad Faisal Cheema, and Gerik Scheuermann. 2015. On Close and Distant Reading in Digital Humanities: A Survey and Future Challenges. *EuroVis (STARs)* 2015 (2015), 83–103.

[7] Alexander Kachkaev, Daniel Wolff, Mathieu Barthet, Dan Tidhar, MD Plumbley, Jason Dykes, and Tillman Weyde. 2014. Visualising chord progressions in music collections: a big data approach. In *9th Conference on Interdisciplinary Musicology - CIM14*.

[8] Vladimir I Levenshtein et al. 1966. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*. Vol. 10. 707–710.

[9] Anna Maria Matuszewska. 2023. Visual presentation and exploration of musical corpora: Case Study: Oskar Kolberg's Opera Omnia. In *Proceedings of the 10th International Conference on Digital Libraries for Musicology*. 82–85.

[10] Franco Moretti. 2000. Conjectures on world literature. *New left review* 2, 1 (2000), 54–68.

[11] Franco Moretti. 2013. *Distant reading*. Verso Books.

[12] Martin Pfleiderer. 2017. The Weimar Jazz Database. *Inside the Jazzomat. New Perspectives for Jazz Research. Schott-Campus, Mainz* (2017).

[13] Martin Pfleiderer, Klaus Frieler, Jakob Abeßer, Wolf-Georg Zaddach, and Benjamin Burkhart. 2017. Inside the Jazzomat. *New Perspectives for Jazz Research* (2017).

[14] João Felipe Pimentel, Leonardo Murta, Vanessa Braganholo, and Juliana Freire. 2019. A large-scale study about quality and reproducibility of jupyter notebooks. In *2019 IEEE/ACM 16th international conference on mining software repositories (MSR)*. IEEE, 507–517.

[15] Nicolas Ruwet and Mark Everist. 1987. Methods of analysis in musicology. *Music analysis* 6, 1/2 (1987), 3–36.

[16] Ivan Shanin and Simon Dixon. 2023. Annotating Jazz Recordings Using Lead Sheet Alignment with Deep Chroma Features. In *2023 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 1–5.

[17] Chris Stover. 2021. Dig That Lick (DTL): Analyzing Large-Scale Data for Melodic Patterns in Jazz Performances. *Journal of the American Musicological Society* 74, 1 (2021), 195–214.

[18] Fangli Su, Yin Zhang, and Zachary Immel. 2020. Digital humanities research: interdisciplinary collaborations, themes and implications to library and information science. *Journal of Documentation* 77, 1 (2020), 143–161.

[19] Ferdian Thung, Tegawende F Bissyande, David Lo, and Lingxiao Jiang. 2013. Network structure of social coding in github. In *2013 17th European conference on software maintenance and reengineering*. IEEE, 323–326.

[20] Tillman Weyde, Stephen Cottrell, Jason Dykes, Emmanouil Benetos, Daniel Wolff, Dan Tidhar, Alexander Kachkaev, Mark Plumbley, Simon Dixon, Mathieu Barthet, et al. 2014. Big data for musicology. In *Proceedings of the 1st International Workshop on Digital Libraries for Musicology*. 1–3.