









Preservation and Promotion of Opera Cultural Heritage: The Experience of La Scala Theatre

Federico Avanzini , Adriano Baratè , Goffredo Haus ,
Luca A. Ludovico  , and Stavros Ntalampiras 

LIM – Laboratorio di Informatica Musicale, Dipartimento di Informatica
“Giovanni Degli Antoni”, Università degli Studi di Milano,
Via G. Celoria 18, 20133 Milano, Italy
{federico.avanzini,adriano.barate,goffredo.haus,luca.ludovico,
stavros.ntalampiras}@unimi.it

Abstract. This paper focuses on music and music-related cultural heritage typically preserved by opera houses, starting from the experience achieved during the long-lasting collaboration between La Scala theater and the Laboratory of Music Informatics of the University of Milan. First, we will mention the most significant results achieved by the project in the fields of preservation, information retrieval and dissemination of cultural heritage through computer-based approaches. Moreover, we will discuss the possibilities offered by new technologies applied to the conservative context of an opera house, including: the multi-layer representation of music information to foster the accessibility of musical content also by non-experts; the adoption of 5G networks to deliver spherical videos of live events, thus opening new scenarios for cultural heritage enjoyment and dissemination; deep learning approaches both to improve internal processes (e.g., back-office applications for music information retrieval) and to offer advanced services to users (e.g., highly-customized experiences).

Keywords: Cultural heritage · Digital technologies · Machine learning · Music · Opera · Preservation · Promotion

1 Introduction

The impact of digital technologies on the preservation, restoration, and fruition of tangible and intangible cultural heritage is indisputable [7]. Music cultural heritage has also been part of this digital revolution.

A particular emphasis has been given to the preservation and fruition of audio documents, which poses non-trivial challenges, as these document require several layers of contextual information to be preserved along with the audio signal. Preservation of audio documents can be categorized into *passive* preservation, meant to protect the original documents from external agents without

alterations, and *active* preservation, which involves the data transfer from the analogue to the digital domain. Despite initial concerns about the use of digital recordings and digital storage media for long-term preservation, the traditional “preserve the original” paradigm has progressively shifted to the “distribution is preservation” idea of digitizing the content and making it available in digital libraries [11].

The categories of passive and active preservation may be extended to the field of physical artifacts, e.g., musical instruments [2]. Here active preservation involves virtual simulations enabling new means of interaction with objects that are otherwise not accessible, with important implications for museum exhibits in particular. Presenting artifacts to the general public is a complex task: interactive museum installations can increase the engagement and participation of visitors [23], and enforce new forms of learning. Ultimately, applying new technologies to interactive museum installations can create stronger consensus and interest for the preservation of cultural heritage [27].

This paper presents the results of a long-standing project [15] on a specific case study, the historical archive of the La Scala theatre of Milan (*Teatro alla Scala di Milano*), one of the best known musical temples of the world. This is a particularly interesting and challenging case, due to the richness and the heterogeneity of the documents and artifacts that are typically found in an opera archive. The results achieved by the project so far provide the basis for reflecting on the added values of digital technologies applied to the preservation and exploitation of opera cultural heritage, and for proposing further developments building on current research advancements.

The rest of the paper is structured as follows: Sect. 2 will describe the background of the project, Sect. 3 will report the main results achieved by the 10-years cooperation between La Scala and the LIM, Sect. 4 will present a number of scenarios where digital technologies can help the preservation and exploitation of opera cultural heritage, Sect. 5 will focus on machine learning as a promising technology to give new value to archive materials, and Sect. 6 will draw conclusions.

2 Background

In its 35 years of activity, the *Laboratorio di Informatica Musicale* (LIM) of the Department of Computer Science, University of Milan has established collaborations with relevant institutions active in the fields of music production, research, and cultural heritage. Examples include, to cite but a few, Bolshoi Theatre (Moscow), Ricordi Historical Archive (Milan), Bach Archiv (Leipzig), and Paul Sacher Stiftung (Basel).

A particularly important partnership was the one with La Scala theatre of Milan. Such a cooperation lasted more than one decade (from 1996 on) and involved different competences and professional skills available in the LIM staff, ranging from researchers dealing with sound and music computing to students enrolled in the course of Music Informatics.

The performing activities of La Scala embrace operas, concerts, recitals, and a number of other cultural events such as presentations and conferences. Along with production, this theater has also the goal of preserving its cultural heritage, distributed across different organizational structures: the Music Archive (hosting both scores and audio recordings), the Photo/Video Archive, the Sketch Archive, the Costume Archive, the Set Design Archive, and the Properties Archive.

When the joint project started, the ICT revolution was in an early phase.¹ As a consequence, most of the information contained in the archives was either on paper or stored in magnetic tapes. In 1997, a pool of international sponsors (Accenture, HP, Oracle, TDK, etc.) funded a project whose ultimate goal was recovering, restoring and preserving the musical heritage of La Scala theater, through a comprehensive digitization campaign and the creation of an on-line database for their retrieval.

In this framework, the cooperation between the LIM and La Scala mainly achieved three results:

1. The preservation and restoration of the historical audio archive, counting about 5000 open-reel magnetic tapes recorded since 1950;
2. The release of advanced computer-based solutions to foster quick and effective access to information. This part of the project included both a comprehensive relational database to federate heterogeneous contents coming from different archives and an experimental platform aimed at providing multimodal access to information (e.g., via queries by humming, symbolic inputs, etc.);
3. The design and implementation of an integrated asset management system, called *LaScalaDAM*, embracing all the archives of the theatre. The idea was to provide full network access via intranet to La Scala authorized staff, and a filtered view (e.g., only selected items, at lower resolution and higher compression, with watermarking, etc.) to the web users of a dedicated portal.

Even if focusing on different aspects – i.e. digitization, information retrieval, and dissemination – all these experiences share a common goal, namely the preservation and promotion of the rich and heterogeneous cultural heritage owned by an opera house or a similar institution. The mentioned activities will be described in detail in Sect. 3.

3 Achieved Results

This section focuses on the state of the art of the collaboration between La Scala and the LIM, reporting the most relevant results achieved in the field of opera-related cultural heritage.

¹ If, on the one hand, the CD-DA standard for audio was available since 1980, on the other the birth of the World Wide Web dated back to just 5 years before, and digital cameras were not widely adopted in a professional context.

3.1 Digitizing with the Aim to Preserve

The digitization campaign began in 1996 with a rescuing project addressing the audio archive. This first step involved about 5000 open-reel magnetic tapes containing opera, ballet and symphonic-music recordings from 1950 onward. To this end, the LIM set up a digitization environment in a dedicated space at the Department of Computer Science (see Fig. 1).

In detail, a staff of experts prepared the tapes and – when needed – restored them through thermal pre-treatment (see Fig. 2), acquired audio content through professional devices, converted and transferred them to digital media, and collected and organized the corresponding metadata. Thanks to this joint effort, about 10000 h of music were recovered and preserved. Such recordings include performances by world-renowned singers (e.g., Maria Callas, Giuseppe Di Stefano, Mario Del Monaco, Plácido Domingo, Mirella Freni, Luciano Pavarotti, Joan Sutherland, Renata Tebaldi, etc.), great conductors (e.g., Claudio Abbado, Daniel Barenboim, Pierre Boulez, Victor De Sabata, Carlo Maria Giulini, Herbert von Karajan, Carlos Kleiber, Riccardo Muti, Arturo Toscanini, etc.), and famous dancers and choreographers (e.g., Roberto Bolle, Alessandra Ferri, Carla Fracci, Rudolf Nureyev, Roland Petit, etc.).

Since 1998, the digitization project was incrementally extended to all theater's archives, namely the costumes and accessories warehouse, the sketches and costume-designs archive, the props warehouse, the photos and posters archive. Currently, the digitized heritage of La Scala embraces about 24000 sketches and costume designs, 45000 costumes, 60000 accessories (including jewels, clothing, footwear, wigs and hats), and 80000 props. Moreover, the theatre's activity is documented in 17000 posters and in more than one million photographs taken from the stage, rehearsals and back-stage.² For each category, great artists are involved; to cite but a few: Alberto Burri, Dino Buzzati, Salvatore Fiume, Renato Guttuso, and Mario Sironi as sketch and costume designers; Caramba, Emanuele Luzzati, and Franca Squarciapino as fashion designers; Alessandro and Nicola Benois, Ezio Frigerio, Pier Luigi Pizzi, and Robert Wilson as set designers and directors.

The number of digitized objects, the heterogeneity of the acquired materials, and the stature of the artists involved make the cultural implications of this activity evident. This is even more true since the long-term goal was not only to preserve digital objects and their metadata, but also to organize, analyze and disseminate them, as explained in next sections.

3.2 Analyzing and Structuring

Following the chronological order, in this section we will describe the two main initiatives that characterized the collaboration between La Scala and the LIM concerning the access to digital cultural heritage.

A first step, experimental and highly innovative at that time, was represented by the *Music Archive Information System* (MAIS) initiative [12]. The goal of

² Source: <http://www.teatroallascala.org/en/archive/the-historical-archive.html>.



Fig. 1. Devices for open-reel magnetic tapes in the LIM digitization studio: from left to right, Revox A77, Revox B77 MKII and Otari MX-55.

MAIS was to allow content-based queries on available scores and audio materials, so as to support artists in the preparation of a performance already present in the archive. In detail, it was possible to retrieve all the scores and/or audio which contained an exact sequence of notes or a similar one.

From a technical point of view, MAIS was a hardware, software and network environment based on an Oracle 8 object-relational technology. Such a solution was multi-platform, running under Unix, Windows NT and Macintosh operating systems, and distributed, supporting 10 workstations. The environment included also a CD-audio jukebox able to contain up to 224 compact discs and to provide quick access to their content via network (see Fig. 3).

The MAIS overall architecture is shown in Fig. 4. The database was organized around single events (i.e., performances), and stored and linked related meta-data, audio materials (CD-DAs and tapes), and music scores. Both traditional and audio-based queries could be used to retrieve information. Designing MAIS posed several challenges, ranging from large storage requirements to the intrinsic complexity of the database, being the scenario characterized by a large amount of music and multimedia unstructured data stored in a variety of formats.

If on one side the MAIS project was an experimental environment addressing the theatre's staff, a later initiative, called *La Scala DAM* (Digital Asset Management), targeted a wider audience. This project, developed in cooperation with the LIM and other technological partners, equipped the theater with an integrated management system to access its huge digital heritage, from the



Fig. 2. The Heraeus UT6200 oven used to thermally pretreat magnetic tapes.



Fig. 3. The CD-R jukebox in use in the MAIS project.

second decade of the 20th Century to present. In 2006, La Scala DAM was connected to the theater's intranet, thus supporting the everyday activities of every sector and constituting an internal authoritative source for the creation, production and documentation of each show. Based on a relational database, whose structure is described elsewhere [14], La Scala DAM supplied integrated views on performances that included metadata, sketches and costume designs, costumes, footwear, jewellery, head-dress, props, fliers and posters, photographs and audio recordings. A key page of the intranet application is shown in Fig. 5, that provided a synoptic view of events with easily accessible links towards all related metadata and materials. A second example is illustrated in Fig. 6: from this view focusing on photographic material, it was possible to jump to detailed photo descriptions, to enlarge images and to follow links towards other sections (e.g., all the materials related to a given opera, all the titles in that season, all the pictures taken in a date or by a photographer, etc.).

3.3 Disseminating

The cooperation between La Scala and the LIM paved also the way toward the public availability and dissemination of preserved cultural heritage. In fact, an evolution of La Scala DAM project mentioned in Sect. 3.2 was the opening of the historical-archive section, called *ArchivioLaScala*, within the institutional web site of the theater.³ Even if limited and simplified with regard to the huge quantity of information contained in the complete archive, *ArchivioLaScala* represents the web-accessible page of La Scala DAM. The initiative addresses a worldwide audience of scholars and enthusiasts who now have the opportunity to access a relevant part of the cultural heritage of La Scala.

³ <http://www.teatroallascala.org/archivio/>.

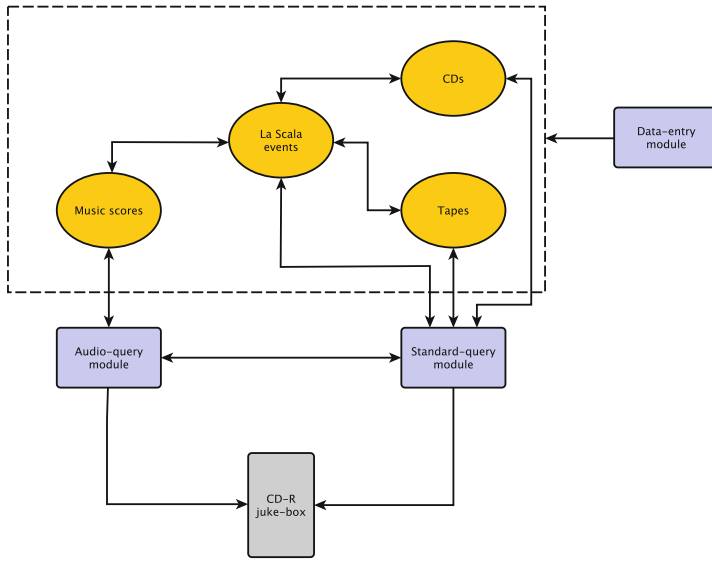


Fig. 4. The MAIS architecture.

This kind of approaches has many practical consequences. The most obvious one is to open cultural heritage archives usually accessible to experts and insiders only. In addition to satisfying the curiosity of regular opera attendees, a long-term goal concerns the involvement of a new audience. At La Scala this aim is pursued by showcasing materials from the past together with teasers about current productions.

Moreover, a suitable combination of digital materials and advanced technologies can be the first step towards information sharing for the constitution of semantic networks capable to offer new services and improved user experience [5]. Such an approach implies the concept of *openness*, in terms of accessibility of knowledge, technology and digital resources. Unfortunately, cultural heritage stakeholders often do not recognize the advantages offered by a shared access to information, rather fearing the infringement of intellectual property and, more in general, the loss of control over their own assets.

Conversely, once the digitization chain is established and tested, benefits may come at a very little cost for the institution. For example, the cooperation between the LIM and Archivio Storico Ricordi, aiming once again at the preservation and promotion of music cultural heritage, has resulted in a number of significant and profitable projects, including the publishing of high-quality reprints of manuscripts and the organization of technology-empowered exhibitions around the world.

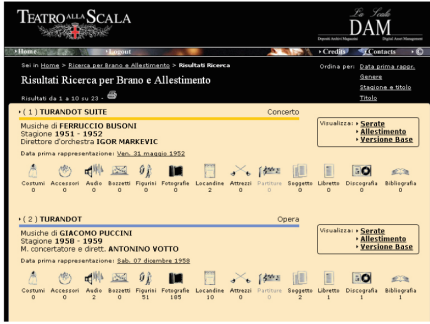


Fig. 5. The synoptic view of events.

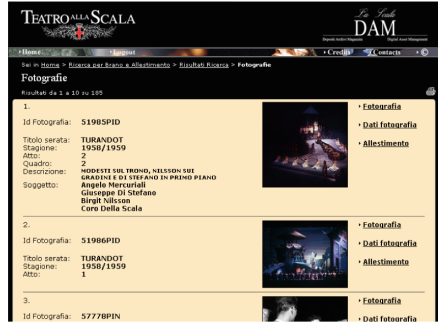


Fig. 6. A result page for photos.

4 Added Value of Opera Cultural Heritage

From the experience acquired working with La Scala and other musical institutions, it is possible to infer the advantages that an opera house can expect from the adoption of ICT approaches.

First, it is fundamental to preserve the past, also with the aim to provide new value to cultural heritage and exploit archive materials. In this sense, digitization campaigns conducted through professional devices and following international standards are only the first step to take. For example, an underestimated aspect is the need to monitor continuously the conservation status of media: contrary to what common sense suggests, “digital” does not mean “everlasting”. As pointed out in [9], preserving digital information is plagued by short media life, obsolete hardware and software, slow read times of old media, and defunct standards.

The availability of cultural heritage in digital format can open previously unexplored commercial opportunities to opera houses, going beyond traditional editorial initiatives and merchandising. For example, new media products can be conceived, containing a number of historical performances of the same opera, all synchronized with different versions of the score. This scenario is made possible by a multi-layer environment (i.e., a format plus an application) for the description of music information, such as the one mentioned in [4].

One further goal of a modern opera house is to enlarge its audience, and, in this perspective, young spectators represent the future. Meeting the expectations of digital natives implies implementing novel ways to let them experience the show and offering advanced services. Moreover, the use of suitable techniques to record live performances and the adoption of new network technologies to broadcast them can extend the potential audience by overcoming the limits of space and time distance. An example is *augmented opera* [5], i.e. the application of augmented reality to opera live experience. An ad-hoc application installed on users’ devices could present features such as automatic score and libretto following with multi-language translation, dynamic identification of characters, links to user-tailored content and external services, etc. In this way, the ongoing show is enriched by already-available archive materials (e.g., scores, librettos, on-stage

photos from the past, stage maps, etc.) in order to offer a more comprehensive and engaging user experience.

A technology that in the near future will foster the promotion of opera cultural heritage, embracing both new productions and archive materials, is the one of 5G networks. Thanks to the expected characteristics of 5G – namely improved bandwidth, reliability, and density of devices in an area – it will be possible to release innovative services for mobile devices rich in multimedia content, supporting multimodal interaction, and highly customizable depending on users’ requirements and special needs. The applications are manifold, ranging from multi-angle and multi-track experience of opera to immersive navigation inside performances [3]. Materials from the past can be integrated with ad-hoc recordings and new digital objects within a unique environment, for both entertainment and educational purposes [3].

It is worth emphasizing that, nowadays, a considerable problem for opera houses is the economic sustainability of their activities. For example, La Scala is still far from the goal of financial self-sufficiency, rather relying on significant public and private funding to achieve a balanced budget. In this sense, both the promotion of the preserved cultural heritage and the release of new services and products based (also) on archive materials can be a low-cost solution to generate income from already available assets.

5 Future Perspectives: Machine Learning for MIR Purposes

In this section, we will focus on the application of machine learning to music information retrieval (MIR) for opera cultural heritage, since it can represent a promising approach both for scientific investigation (e.g., musicological analysis) and for practical aims (e.g., multilayer content navigation and retrieval).

In this context, technologies can be employed to automatically process the above-mentioned data sources in either a mono- or multi-modal way following the line of thought described in [24]. In such a context, the MIR framework can elaborate either on a single data source alone or on multiple ones at the same time and facilitate several applications. With the noticeable exception of the experimentation conducted in the MAIS project, indexing, retrieval and navigation within the La Scala dataset is being carried out in a manual fashion so far without efficient description of music content. Indexing and retrieval services can automatically return information related to queries such as “locate measure x in all the available recordings of a given piece”, “retrieve the performances more similar to the one of a given soloist/conductor”, “find all a-capella music pieces”, etc. Such queries need an internal characterization of music content enabling tagging information related to melody, rhythm, harmony, expression, etc.

During the last decades, MIR has been receiving ever-increasing attention including on-line music services able to reply to such queries via interpretation of musical content. Recognition of musical content, similar to generalized sound

recognition, is based on the fact that each music piece has a characteristic way of distributing its energy on its frequency content, which constitutes its so-called signature [21]. MIR technologies offer automatic extraction of information by processing the available audio signals as the main modality as well as complementary ones such as music scores, video of the performances, etc. Thus, it paved the way for novel applications in the context of an opera house, such as:

- *Music genre classification*—Genre is one of the most widely-used descriptions of music. The literature typically includes genres such as classical, reggae, jazz, rock, etc. [18], but in the case of La Scala archive we can include a more detailed characterization in terms of historical period, style, cultural influences, etc. Such a categorization is subjective up to a certain extent, but in this dataset we looked for overall agreement between experts of this type of music. In general, music signals belonging to the same category share several properties (instrumentation, rhythmic and harmonic structure, etc.) allowing us to assume that each category groups pieces that sound similarly. Features of particular interest here try to capture information associated to instrumentation, rhythm, pitch, harmony, timbre, etc. [26].
- *Music transcription*—This specific application involves mapping the audio signals of the La Scala archive to its complete score [6]. Such algorithms typically estimate the pitch and form the symbols accordingly. However, there are several obstacles here, such as multiple notes emitted by the same instrument or orchestral group, temporal overlapping of several instruments, etc. For example, a reliable algorithm for music transcription could automatically obtain the score of an improvisation (e.g., the extemporaneous *cadenza* by a soloist), the underlying symbolic representation in queries by humming, the audio-driven transcription of non-traditional graphical scores into common western notation.
- *Performer recognition*—It is widely accepted that great soloists and conductors are characterized by their own performance style. Even though there is no formal and explicit definition of what style is, we can trace it back to expression-related parameters like variations in dynamics (loudness) and agogics (tempo), characterization of certain notes/passages, etc. This type of information is not included in the music score, so its investigation is particularly interesting. Despite the subjectivity of the task, there are many researchers who have employed low-level features and machine learning algorithms to identify performers [10].
- *Instrument recognition*—This application encompasses the recognition of both type and number of instruments included in a music piece [13, 17]. Such a high-level description comprises a useful information towards improved organization/searching of the La Scala archive, since the ensemble can characterize a piece, genre, historical period, etc. At the same time, such information can boost the signal processing front-end facilitating multi-pitch analysis, source separation, synchronization, and transcription [21].
- *Music emotion recognition*—This application domain is concerned with the prediction of the emotion perceived by the audience of a specific music piece.

Music, as a means of communication, evokes certain emotional responses. Identification of the emotion on the listener's side may be indicative of the respective human reaction. To this end, appropriate signal processing and pattern recognition algorithms can be employed [19].

This is a non-exhaustive list of scenarios, as the richness of the La Scala archives is still under exploration and emerging applications might arise targeting the particular character of such a cultural-defining theater. Interestingly, all the above-mentioned applications as well as emerging ones can be viewed by means of a chronological prism towards a historical analysis of various aspects including how specific performers, performances, compositions, etc. have been evolved with the passage of time.

As a general comment, feature extraction is typically carried out on small parts of audio data (frames) where we assume that the signal holds stationary properties. This assumption is not restrictive as vocal articulators, musical instruments, singers, etc. have an inertia that prohibits instantaneous changes in the frequency content of the sound. As mentioned above, there are several features serving various applications while a rather exhaustive summary is offered in [1]. Interestingly, during the last years, following the boom of Deep Learning, several researchers have tried to address MIR-related problems in an end-to-end fashion, i.e. without any preprocessing/feature extraction step utilizing the audio signal alone with promising results [20]. Such approaches remove the need of designing handcrafted features based on domain knowledge since they elaborate directly on the time and/or frequency domain [25] with the ability to automatically characterize musical content [22].

Finally, deep networks allow generation of musical content as well [8, 16], which could be a fruitful direction considering the relevance of La Scala performances.

6 Conclusion

In this paper we have summarized the activities a long-standing project on the historical archive of the La Scala theatre of Milan, emphasizing the most interesting and challenging aspects of this specific case study, and focusing on the main results achieved so far. Based on these, we have proposed a reflection on the implications of this work, in terms of improved access to this heritage, augmented experiences of operas, and novel commercial opportunities. Finally, we have discussed a particularly promising direction of research, which leverages on recent developments in the field of machine learning to support a variety of music information retrieval tasks that can be fruitfully applied to several application scenarios and open up new forms of interaction with opera cultural heritage.

References

1. Alías, F., Socoró, J., Sevillano, X.: A review of physical and perceptual feature extraction techniques for speech, music and environmental sounds. *Appl. Sci.* **6**(5), 143 (2016). <https://doi.org/10.3390/app6050143>

2. Avanzini, F., et al.: Virtual reconstruction of an ancient Greek pan flute. In: Proceedings of International Conference Sound and Music Computing (SMC 2016), Hamburg, pp. 41–46 (2016)
3. Baratè, A., Haus, G., Ludovico, L.A., Pagani, E., Scarabottolo, N.: 5G technology and its applications to music education. In: Multi Conference on Computer Science and Information Systems, MCCSIS 2019 - Proceedings of the International Conference on e-Learning 2019, pp. 65–72. IADIS Press (2019)
4. Baratè, A., Ludovico, L.A.: IEEE 1599 applications for entertainment and education. In: Baggi, D., Haus, G. (eds.) Music Navigation with Symbols and Layers: Toward Content Browsing with IEEE 1599 XML Encoding, pp. 115–132. Wiley, Hoboken (2013). <https://doi.org/10.1002/9781118494455.ch7>
5. Baratè, A., Ludovico, L.A.: Local and global semantic networks for the representation of music information. *J. e-Learning Knowl. Soc.* **12**(4), 109–123 (2016)
6. Benetos, E., Dixon, S., Duan, Z., Ewert, S.: Automatic music transcription: an overview. *IEEE Signal Process. Mag.* **36**(1), 20–30 (2019). <https://doi.org/10.1109/MSP.2018.2869928>
7. Cameron, F., Kenderdine, S.: *Theorizing Digital Cultural Heritage: A Critical Discourse*. MIT Press, Cambridge (2007)
8. Chemla-Romeu-Santos, A., Ntalampiras, S., Esling, P., Haus, G., Assayag, G.: Cross-modal variational inference for bijective signal-symbol translation. In: Proceedings of the 22nd International Conference on Digital Audio Effects, DAFx 2019, Birmingham, UK, pp. 1–8, September 2019
9. Chen, S.S.: The paradox of digital preservation. *Computer* **34**(3), 24–28 (2001)
10. Chudy, M., Dixon, S.: Towards music performer recognition using timbre. In: Proceedings of the 3rd International Conference of Students of Systematic Musicology, pp. 45–50 (2010)
11. Cohen, E.: Preservation of audio in folk heritage collections in crisis. In: Proceedings of Council on Library and Information Resources, pp. 65–82 (2001)
12. Ferrari, E., Haus, G.: The musical archive information system at Teatro alla Scala. In: Proceedings IEEE International Conference on Multimedia Computing and Systems, vol. 1, pp. 817–821. IEEE (1999)
13. Gururani, S., Sharma, M., Lerch, A.: An attention mechanism for musical instrument recognition. *CoRR* abs/1907.04294 (2019)
14. Haus, G., Ludovico, L.A.: The digital opera house: an architecture for multimedia databases. *J. Cult. Herit.* **7**(2), 92–97 (2006). <https://doi.org/10.1016/j.culher.2006.02.007>
15. Haus, G., Pelegrin Pajuelo, M.L.: Music processing technologies for rescuing music archives at Teatro alla Scala and Bolshoi theatre. *J. New Music Res.* **30**(4), 381–388 (2001)
16. Kumar, K., et al.: Melgan: generative adversarial networks for conditional waveform synthesis. In: Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché Buc, F., Fox, E., Garnett, R. (eds.) *Advances in Neural Information Processing Systems* 32, pp. 14881–14892. Curran Associates, Inc. (2019)
17. Nagawade, M.S., Ratnaparkhe, V.R.: Musical instrument identification using MFCC. In: 2017 2nd IEEE International Conference on Recent Trends in Electronics, Information Communication Technology (RTEICT), pp. 2198–2202, May 2017. <https://doi.org/10.1109/RTEICT.2017.8256990>
18. Ntalampiras, S.: Directed acyclic graphs for content based sound, musical genre, and speech emotion classification. *J. New Music Res.* **43**(2), 173–182 (2014). <https://doi.org/10.1080/09298215.2013.859709>

19. Ntalampiras, S.: A transfer learning framework for predicting the emotional content of generalized sound events. *J. Acoust. Soc. Am.* **141**(3), 1694–1701 (2017). <https://doi.org/10.1121/1.4977749>
20. Pons, J., Nieto, O., Prockup, M., Schmidt, E.M., Ehmann, A.F., Serra, X.: End-to-end learning for music audio tagging at scale. CoRR abs/1711.02520 (2017)
21. Potamitis, I., Ganchev, T.: Generalized recognition of sound events: approaches and applications. In: Tsihrintzis, G.A., Jain, L.C. (eds.) *Multimedia Services in Intelligent Environments*. SCI, vol. 120, pp. 41–79. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78502-6_3
22. Roman, M., Pertusa, A., Calvo-Zaragoza, J.: A holistic approach to polyphonic music transcription with neural networks. In: *Proceedings of the 20th International Society for Music Information Retrieval Conference, ISMIR, Delft, The Netherlands*, pp. 731–737, November 2019. <https://doi.org/10.5281/zenodo.3527914>
23. Simon, N.: *The Participatory Museum*. Museum 2.0 (2010). Creative Commons
24. Simonetta, F., Ntalampiras, S., Avanzini, F.: Multimodal music information processing and retrieval: survey and future challenges. In: *2019 International Workshop on Multilayer Music Representation and Processing (MMRP)*, pp. 10–18, January 2019. <https://doi.org/10.1109/MMRP.2019.00012>
25. Simonetta, F., Chacón, C.E.C., Ntalampiras, S., Widmer, G.: A convolutional approach to melody line identification in symbolic scores. In: *International Society for Music Information Retrieval Conference*, pp.1–8 (2019)
26. Sturm, B.L.: A survey of evaluation in music genre recognition. In: Nürnberg, A., Stober, S., Larsen, B., Detyniecki, M. (eds.) *AMR 2012*. LNCS, vol. 8382, pp. 29–66. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-12093-5_2
27. Styliani, S., Fotis, L., Kostas, K., Petros, P.: Virtual museums, a survey and some issues for consideration. *J. Cult. Herit.* **10**(4), 520–528 (2009)