Perspective

Sharing Cultural Heritage—The Case of the Lodovico Media Library

Matteo Al Kalak * and **Lorenzo Baraldi ***

Research Center on Digital Humanities, University of Modena and Reggio Emilia, 41121 Modena, Italy
* Correspondence: matteo.alkalak@unimore.it (M.A.K.); lorenzo.baraldi@unimore.it (L.B.)

Abstract: The article aims to reflect on the Lodovico media library, a digital repository preserving the digitised cultural heritage of the Emilia-Romagna region. The first part covers the project's history and the challenges encountered during its setup phase, and we also explore the co-creation approach employed in defining the metadata architecture. The discussion extends by outlining the key features of shared metadata, illustrating their application to diverse digital objects within the Lodovico media library. Following a concise examination of the methodology for collecting/creating data and the initial research findings, the article concludes by highlighting the project's potential in the realm of automatic handwriting recognition processes.

1. Introduction

Cultural institutions, both public and private, are becoming increasingly aware of the need to make their heritage available in digital forms and the opportunities that this can generate. The current policies of governments support development in this direction, stimulating trans-disciplines like that of the digital humanities, revitalising the central role of the humanities and rethinking their social and political role at this crucial juncture (cf. Burdick et al., 2012 [1]). Over the following pages, we intend to illustrate the solution offered by the Lodovico media library to the digitisation of museums, libraries and archives in a specific region of Italy, Emilia-Romagna. This article will reflect on the key points of the Lodovico project, beginning with the concrete choices that shaped its conception and set-up. Particularly, we will focus on the following aspects: (1) the 'federated' and multitenant architecture of Lodovico; (2) the structure of the metadata, the methods of visualisation and the collaborative potential of the project; (3) the handwriting recognition experiments connected with digitised heritage with consideration of the questions related to computer science research.

2. A Cross-Institutional and Multitenant Media Library: Birth of a Project

2.1. A Universal Library: The Legacy of Lodovico Antonio Muratori

To understand the first point, i.e., the federated nature of the Lodovico media library, it is important to take account of the design process that led to its creation. A clue about the intended nature of the media library lies in its name: Lodovico.

This is a reference to the intellectual Lodovico Antonio Muratori (1672–1750) who, in the 18th century, participated in the so-called "Republic of Letters", an extensive network of scholars whose communications went beyond religious and political boundaries and, today, inspires numerous projects in the digital sphere (Hotson, Wallnig 2019 [2]; Edelstein et al., 2017 [3]). In his works, often vast in scale, Muratori collected and published materials from Italian and European libraries and archives, constructing a kind of ante litteram universal library (Imbruglia 2012 [4]).

Inspired by this legacy, the media library takes his name and was created to gather heterogeneous materials that can contribute to forming a comprehensive and valuable body

of knowledge. The assumption behind this vision is that, by bringing together different data and heritage, it is possible to expand research and knowledge sharing in a highly fragmented context. This context, over time, has been subject to ruptures, dispersions and organisations, forming distinct cultural entities. For this reason, overcoming institutional barriers and virtually reuniting forms of heritage that have been broken up adds value on at least two fronts. Firstly, it makes specific areas of cultural heritage more visible thanks to the possibility of performing integrated searches (in fact, users are rarely capable or interested in searching for documents on the basis of their location). Secondly, it makes it possible to acquire a more accurate understanding of the historical sedimentation of heritage, highlighting its external relations with a specific conservation body.

2.2. Gathering Cultural Institutions

Therefore, Lodovico was conceived according to a multitenant structure. This means that the archive space in which the digital objects of the media library are stored is organised into 'compartments' called tenants. Individual tenants represent virtual spaces in which images and data connected with each institution are hosted in an independent way and separated from others. This, however, does not limit the user as transversal access modalities are also provided. These spaces, indeed, are not organised according to the topic (objects that have similar content), type (i.e., the same types of documents) or chronology. Rather, they correspond to the partners of Lodovico (that is, cultural institutions) so as to promote the institutional dimension. This choice is in line with those adopted in other digital humanities projects (see Presner, Johanson 2009 [5]: p. 3, on contact with the institutions to generate "new knowledge and new forms of civic engagement").

The front end and interface are also shared: individual cultural institutes agree to integrate with each other, not renouncing their identity but making it part of a larger framework. By entering Lodovico, the user can thus directly and easily access a cultural ecosystem.

This architecture aims to make the project sustainable economically, too. By sharing costs between different cultural institutions and scientifically, by connecting the world of research and the institutions that conserve cultural heritage, it is easier to ensure sustainability in the use of digitised cultural heritage over time.

Therefore, the novelty of the project is not in the metadata methodology adopted (see Section 3) nor in the use of innovative technological standards. Rather, it lies in the logic of sharing technological infrastructures on a large territorial scale and between different museums, archives, and libraries. However, Lodovico does not aspire to be a mere aggregator, but a structure that integrates different cultural heritages at the origin in a very different way compared to other media libraries or digital archives (see some examples in Levenberg et al., 2018 [6]).

2.3. The Pilot Case

Although Lodovico today includes heritage spread all over the region of Emilia-Romagna, the pilot phase involved a smaller area before being redesigned and eventually extended to all cultural institutions in the region, as per the following diagram (Figure 1):

The pilot phase involved the area of the city and province of Modena, which can be regarded as indicative of the "scattered" nature of Italian cultural heritage and, more specifically, that of the Emilia-Romagna region.

From 1598 to 1859, Modena was the capital of a state with a long tradition, the Duchy of Este (Signorotto, Tongiorgi 2018 [7]; Fumagalli, Signorotto 2012 [8]; Folin 2001 [9]; Marini 1979 [10]). At the time of the Unification of Italy (1861), the new kingdom "reorganised" its cultural heritage, and Modena boasted the prestigious art collection of the Estensi and their ancient archive, with documents dating from the Carolingian era onwards. Modena's precious library was also reorganised to contain rare manuscripts, which are still of great importance for the world of research (see, for example, Di Pietro Lombardi 2017 [11]). At the same time, the city had an ancient past, irrespective of the dynasty of

duchies that governed it, as its Romanesque cathedral testified. Modena had been the site of a medieval commune which for centuries maintained and safeguarded forms of administrative autonomy, as documented in the papers conserved in the Historical Archive of the Municipality of Modena (Bonacini 2002 [12]; Borsari 2001 [13]; Biondi 1987 [14]). Earlier still, it had been ruled by the bishops, whose documents, conserved in the Historical Archive of the Archdiocese, date as far back as the 7th–8th century (Bonacini 2001 [15]; Vigarani 2003 [16]). The interdependence between the various sources of documentary heritage is undeniable, and it is not possible to understand the history of this land without bringing together all of the aforementioned elements (cf. Golinelli 2011 [17]).
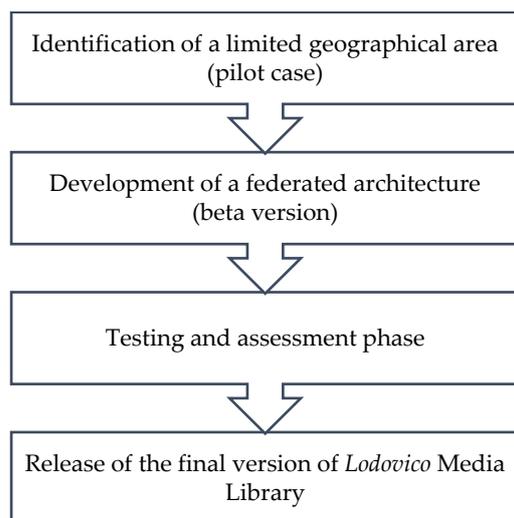

Figure 1. Diagram of the pilot case development.

To involve and consider the requirements of the various cultural institutions, a participatory process was adopted. The first group of institutions interested in the project created an "open" consortium that made allowances for the possible arrival of new partners. Participating parties, coordinated by the Interdepartmental Research Centre on Digital Humanities of the University of Modena and Reggio Emilia (https://www.dhmore.unimore.it/, accessed on 20 November 2023), launched a process for the joint definition of the descriptive standards to apply to the digitised heritage. This co-creation process entailed meetings held to define the needs of the various cultural institutions, identify the specific features of the objects to be digitised and create metadata (meaning the information relating to the digitised objects). Additionally, it involved the development of an initial trial release (beta version) and, following an assessment phase, the release of the final version of Lodovico. The partnership was then expanded to the regional level, including institutions that were not involved in the pilot project. In this regard, it should be noted how the assessment and improvement of Lodovico is an ongoing process driven by the arrival of new contributors to the media library and, therefore, the identification of new needs.

To facilitate the described co-creation process, it was decided to use the Italian language (with which the cultural institutions involved are most familiar). A way of providing the metadata in English is currently under consideration; also in view is the possibility of sharing Lodovico's digital objects on Europeana.

3. Cross-Typological Metadata for a Non-Specialist Public

3.1. Metadata Structure

It is common knowledge that it is advisable to adopt shared metadata standards in order to facilitate interoperability. In addition to this, a "federated" project like Lodovico must clearly develop information standardisation protocols that can be applied by all the institutions participating in the project. In this way, it can be ensured that data are interoperable in terms of structure and content.

To ensure that the metadata are connected, accessible and interoperating, a metadata model (the set of data to describe the objects hosted in Lodovico) was prepared in compliance with the FAIR principles (https://www.go-fair.org/fair-principles, accessed on 20 November 2023) and on the basis of the Dublin Core standard (https://www.dublincore.org/, accessed on 20 November 2023). The categories of the Dublin Core were adapted and matched with the description needs of the cultural institutions as well as with the user/research questions.

A summary is provided below (Table 1):

Table 1. Lodovico Categories Description.

Lodovico Categories (Dublin Core Based)	Description	User/Research Questions
Header	Shelf mark (actual and old shelfmark, if any)	Where is preserved the original object?
Title	Original or attributed title for a short description of the digitised object	What is the digitised object?
Chronological date	Chronological extremes	When was the original object produced?
Physical description	Medium (type and material), dimensions (height, width, depth), and consistency (quantity, unit of measurement) of the digitised object	What are the material characteristics of the original object?
Content description	Summary or transcription of the contents	What is the content of the original object?
Persons	Author, recipient/addressee, people cited, possessor, donor (when persons)	Which persons are related to the original object? How (authors, addressees, cited, etc.)?
Entities	Author, recipient/addressee, entities/institutions cited, conservator (when entities)	Which entities are related to the original object? How (authors, addressees, cited, etc.)?
Places	Topical date, places cited	Where was the original object produced? What other places are mentioned?
Subject	Topic(s) related with the contents of the digitised object	What are the topics with which the original object and its contents can be associated?
Notes	Free-text field for additional information	Is there any other useful or necessary information to be given to the user/researcher?
Language	Language(s) used for the contents	What language is used in the original object?

Without going into detail, we can see that eleven categories of metadata were identified. Each one is associated with one or more specific types of metadata that specify the information that the users can access.

The federated structure of Lodovico did not just bring together different cultural institutions but also involved different types of documents to be digitised. Though representing a very clear choice, one of the most salient features of Lodovico is its aim towards a non-specialist audience, like many other digital projects (cf. Previtali 2022 [18]: pp. 37–39). With this in mind, we decided to avoid different categories of metadata on the basis of document type (archive documents, library items, museum items, geographical maps, musical documents, etc.). Therefore, all of the digital objects described in Lodovico use the same descriptive categories in order to maximise the possibility of creating a dialogue between the objects, regardless of their nature.

As mentioned, this was not an easy decision. Indeed, it required participating institutions to change the language they were used to. Also, the desire to engage a non-specialist public via a single "set of grammar rules" meant introducing rationalisations whose repercussions in terms of the knowledge and correct understanding of the data are constantly monitored within the project.

Nevertheless, in an era in which digital technology has profoundly changed traditional practices like archiving (cf. Ernst 2012 [19]), programmatically speaking, Lodovico

represents a medium for sharing cultural heritage on a vast scale. In relation to a special-ist public—which is not the primary target of the project—it can, therefore, be viewed as a tool for tapping into more complex knowledge and digitally accessing a legacy of information that the expert user is able to decipher in depth. This approach is even more valuable when, as is the case with Lodovico, institutions which have used highly specialist tools for cataloguing their heritage are involved. For example, many archives have constructed descriptive trees using software like x-Dams (https://www.xdams.org/, accessed on 20 November 2023) or, in the case of ecclesiastical archives, CEI-Ar (https://www.beweb.chiesacattolica.it/beniarchivistici/aggregatore/30/Il+progetto+CEI-Ar, accessed on 20 November 2023), which represent useful sources of supplementary information to that made accessible by Lodovico.

3.2. A User-Friendly Medialibrary

Linked to this is the way in which the metadata and its architecture are presented visually to the user (Figure 2), which is a choice that every digital project is required to make, with the acknowledgement of the fact that no representation is neutral (cf. Burgio 2021 [20] with regard to infographics).
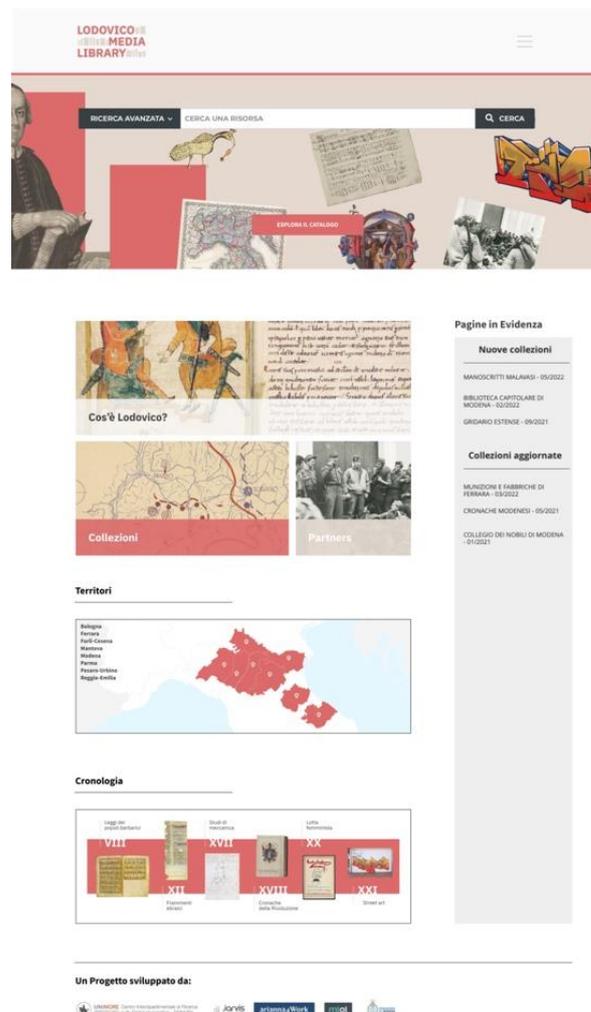

Figure 2. Lodovico homepage (https://lodovico.medialibrary.it/home/index.aspx, accessed on 20 November 2023).

The metadata related to one of the 11 categories listed above is usually presented grouped together or alongside each other to make it easier to read and understand. It is organised into three blocks. Firstly, there is a heading with the key data: a miniature of the

digital object, title, author and date of the document, as well as its type. Secondly, there is the image of the digitised object codified according to the standards defined by the IIIF protocol, which permits its interoperability with other digital libraries (https://iiif.io/, accessed on 20 November 2023). The quality and resolution of the images are designed to pique the interest of the user who, via the viewer, can enlarge the digital object and observe it in a level of detail that, in many cases, would not be possible through a traditional consultation of the materials.

Finally, the descriptive details of the digitized object are reported below the image: an internal description is provided, which may consist of a summary or a more or less analytical indexation of the content, supplemented, where relevant, by a comprehensive transcription. These are followed by data on the location of the document (in an archive, library, museum or other conservation entity), the conservator institution, the people or bodies cited, the places cited, the subject index, the exact chronological date (dd-mm-yyyy), and descriptive details of the medium.

The 11 macrocategories and/or relative metadata do not have to be produced in full for every series or document collection. Indeed, they can be activated to a variable degree according to the cost–benefit ratio and, on a scientific level, on the basis of their actual usefulness. For instance, the printer is an essential item of data for describing a book but not for other printed materials like decrees or brand manifestos; in manuscripts, it is absent. What is essential, on the other hand, is that the same types of information converge in the same categories of metadata and that the latter are codified on the basis of standard criteria in accordance with a compilation manual used by the Lodovico consortium.
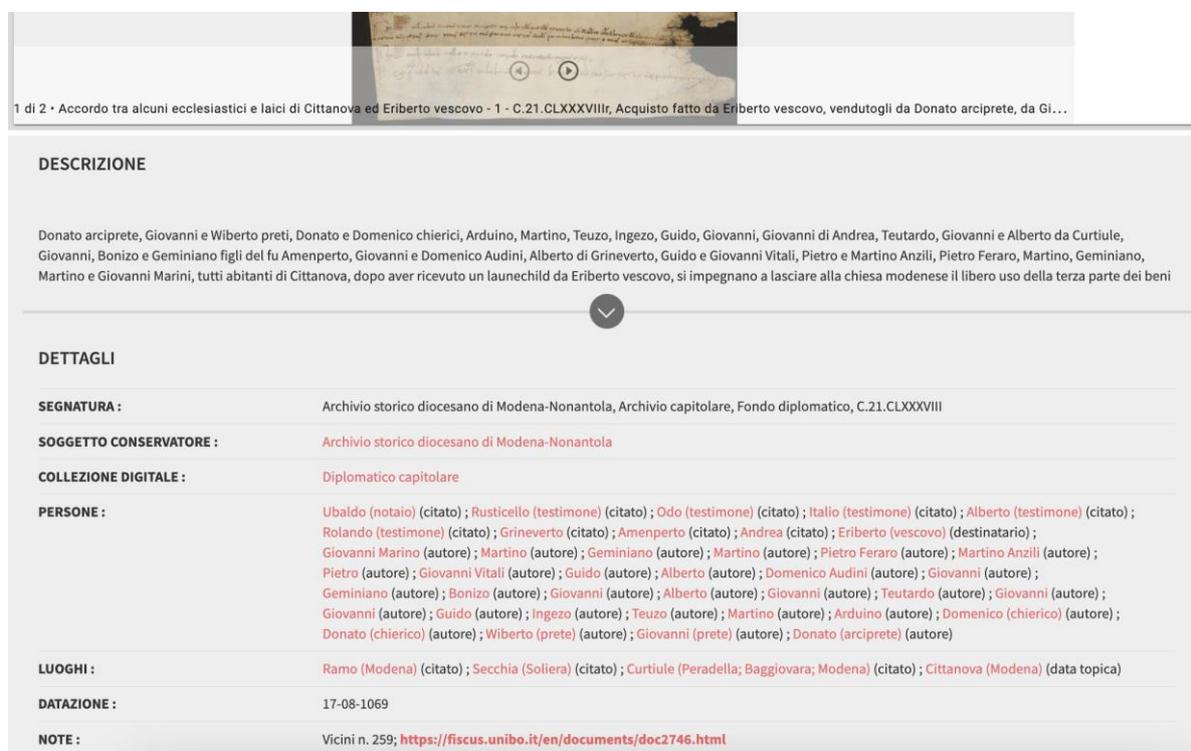
All fields can be searched by users, who can take advantage of a google-like search tool to find information relatively quickly. This, however, does not exclude the possibility of performing advanced searches using a specific query mask that presents the most relevant categories of metadata.

One last useful note: the federated logic of Lodovico and the efforts to guarantee the interoperability of both the metadata and the images ensure that the media library can easily be associated with other projects, according to the inclusive and collaborative spirit typical of the digital humanities (cf. Spiro 2012 [21]). In some cases, these projects have similar characteristics; in many others, they have different aims and objectives.

Just a couple of further examples: On Lodovico, it is now possible to view some of the materials implemented by the Digit.a.re project, which aims to create a digital archive of the photography and graphics collections of Reggio Emilia (https://eventi.comune.re.it/eventi/evento/digit-a-re-digital-archives-reggio-emilia/, accessed on 20 November 2023). Despite being conceived with different architectures and methods of accessibility, Lodovico and Digit.a.re are two digital libraries that adopt the IIIF protocol and gather items targeted at a broad audience. These characteristics have made it possible to share data (including images, which are replicated via an IIIF manifest): the advantage for the Digit.a.re project lies in the possibility of making its heritage (relating exclusively to the materials conserved at the Biblioteca Panizzi of Reggio Emilia) available in a federated environment, forming additional connections and broadening the visibility of the items themselves; at the same time, Lodovico is able to enrich its digital archive and offer users a better service.

A similar logic is employed to enable collaborations with overtly specialist projects. This is the case with the *Fiscus* project, which aims to study the economic and patrimonial dimensions of public powers in Italy in the Late Middle Ages (https://fiscus.unibo.it/en/, accessed on 20 November 2023). *Fiscus* has created a specific database, in part furnished with digitised documents: some of these—scrolls relating to the geographical area of interest of Lodovico—have been "shared". The highly detailed and analytical data on *Fiscus* has been simplified and condensed for its export onto Lodovico, which, in turn, in the field reserved for notes, is able to provide a link back to the descriptive entries of *Fiscus* (Figure 3). Once again, this reciprocity derives not from the similarity of the projects but their complementary nature: *Fiscus* is a specialist project with data deriving from in-depth research and analysis; Lodovico, meanwhile, aims to provide simple information that

facilitates interconnections. The sharing of information and reciprocal references via urls enriches the two projects and, above all, enables them to reach different audiences: *Fiscus* becomes visible to the general public, while Lodovico reaches a highly distinct target with specific demands.

Figure 3. A document on Lodovico shared with the *Fiscus* project.

As is tacitly implied in the methodology of digital projects, Lodovico remains a work in progress in which experimentation in the field helps to refine the theoretical aspects through mistakes, revisions and reconsiderations (cf. Mahomy, Pierazzo 2012 [22]).

3.3. Collecting and Producing Data

Finally, it is necessary to briefly consider the methodology used in relation to metadata production and collection. The project, as mentioned, is promoted and coordinated by DHMoRe, a research centre specialised in digitisation and the metadata of material cultural heritage. Metadata are collected from existing scientific catalogues or provided by cultural heritage institutions from their internal databases. DHMoRe researchers (over 70 junior and senior researchers) check the quality of the data, their completeness and correct standardisation. When there are no cataloguing tools (or these tools are not sufficient or reliable), DHMoRe researchers directly produce the necessary metadata through fieldwork. The entire process of data collection, validation and publication is overseen by a project manager and a scientific PI, who, depending on the project, is supported by other senior researchers with expertise in the various domains (history, literature, philosophy, law, etc.).

This workflow poses a number of challenges: first of all, it seeks to establish a dialogue between the world of academic research and cultural institutions (two worlds which, in the Italian context, are not always integrated and often use different languages). Moreover, the federative model imposes a multidisciplinary approach and forces experts from different fields to use the same data architecture. Lodovico's team is therefore called upon, from time to time, to verify how the general model can/should be adapted to the description of a specific digital object (describing a book is very different to describing a collection of photos or a picture).

This is why, since the testing phase (pilot case), the data architecture has been tested on different types of objects. Similarly, the process of co-creation of the metadata model involved experts from various disciplines and representatives of museums, archives and libraries.

The first results obtained from the research are encouraging: the federative structure described above was, in fact, capable of linking cultural heritage conserved in different institutions, and it was possible to reconstruct unforeseen and, in some cases, unknown connections between documents and objects conserved by Lodovico project partners (for example, thanks to digital cataloguing, the provenance archives of many of the 100,000 autographs conserved in the Autografoteca Campori of the Estense Library have finally been identified; see Al Kalak, Fumagalli 2022 [23]).

4. New Frontiers of Experimentation

In pursuit of a modern and efficient digital library (cf. Baraldi 2018 [24]), the Lodovico project undertakes extensive research in the areas of annotation, analysis, and automatic recognition of text, specifically focusing on the challenging task of Handwritten Text Recognition (HTR) for non-traditional digitised historical handwritten documents with complex layouts. In short, the HTR task requires an algorithm to automatically understand the content of a handwritten document by providing a natural language transcription of its textual content. The availability of an HTR component in a digital library, therefore, aims to facilitate the organisation of knowledge within historical and artistic archives and facilitate the retrieval of relevant information from textual queries without requiring the user to open a manuscript and read it to grasp its content. In the following, we will discuss the steps carried out to develop HTR algorithms for the Lodovico library, describe the data annotation strategy and the HTR algorithms developed, and discuss their efficacy in quantitative terms.

4.1. Analysis of the Visual Characteristics

As a first step towards the development of recognition algorithms, an analysis of the digitised documents available on Lodovico has been conducted to determine their quantity, existing annotations, and visual characteristics. The examination of illuminated manuscripts, particularly the *Codici Muratoriani*, revealed that the paper support exhibits signs of aging, imperfections inherent to the medium (such as ink bleed-through), and significant variation in writing styles, even by the same author. These characteristics are commonly observed in historical documents (Figure 4).

It is evident that the paper support has stains (caused by ink acidification, ink drops, and humidity). It is also noticeable that the text from the underlying page is visible and that there are erasures, corrections, and marginal annotations. The layout of the pages varies, as does the handwriting. All these considerations were then employed as the basis for developing effective HTR algorithms, which will be discussed in the following sections.

Figure 4. Examples of pages taken from the same bundle of the *Codici Muratoriani*.

4.2. Data Annotation Strategy

As all learnable algorithms require training on a significant amount of manually annotated data, a semi-automatic data annotation strategy was planned, which also led to the development of a dedicated annotation platform, "Transcribe," aimed at creating datasets for HTR, containing documents from the corpora preserved in the Modena archives, with which it is directly integrated (such as the Estense Digital Library).

This activity has resulted in the creation of two line-level HTR datasets (i.e., operated line by line): "Leopardi," and "LAM" (cf. Cascianelli et al., 2021 [25]). The first dataset contains autograph letters by Giacomo Leopardi and reflects the typical challenges of HTR on small single-author collections, typically found in historical archives: the collections are of small size, the papers are written on paper support with ink that has been preserved over time in a peculiar way, and the language and style are specific to the author and the era in which they lived. The second dataset, considerably larger (currently the largest line-level HTR dataset), contains autograph letters by Lodovico Antonio Muratori, written over a long period of time, on various types of support, with handwriting that has changed over time and under varied preservation conditions.

4.3. A Short Background on Handwritten Text Recognition

As mentioned, HTR concerns the development of algorithms that can automatically translate and input handwritten text in natural language. Before diving into the HTR algorithms developed specifically for the Lodovico project, we will give the reader a short background on the state of the art of HTR algorithm development. Generally speaking, HTR can be tackled by considering different textual elements, i.e., characters, words, lines, paragraphs, or pages. Line-level HTR is among the most popular variants and can be performed on pre-segmented text or used in combination with layout analysis and line-level segmentation to obtain a paragraph-level or page-level HTR system.

Back in the origins, HTR was tackled by applying simple models like Hidden Markov Models for image representation and n-gram-based language models for textual output prediction (Toselli et al., 2004 [26]). Later, when the deep learning revolution kicked in, researchers started to employ deep learning for HTR as well. The first deep-learning-based solution to HTR was proposed in (Graves et al., 2008 [27]), where multi-dimensional Long Short-Term Memory networks (MDLSTM-RNNs) are used to build a 2D representation of

the textual image, which is then collapsed into a sequence of vectors used for decoding the output sentence. On this line, a recent trend (Zhang et al., 2019 [28]) entails treating HTR as a sequence-to-sequence problem, which goes from a sequence of text image slices to a sequence of transcribed text generated using a separate recurrent block.

It is also worth saying a few words about OCR (Optical Character Recognition) and its relationship to HTR. As is well known, OCR is the task of transcribing printed text into natural language. Compared to OCR, HTR instead deals with handwritten text. Therefore, it faces challenges related to the high variability of characters in terms of shape and size. A common strategy to tackle this issue is performing specific data augmentation and preprocessing (Wigington et al., 2017 [29]); however, few works have faced this issue at the architectural design level. For example, in (Zhong et al., 2016 [30]), a Spatial Transformer Network was employed for character-level HTR, while in (Bhunia et al., 2019 [31]), an adversarial deformation module was used to warp intermediate convolutional features in a word-level HTR model to help the network deal with character-level variations.

4.4. HTR Algorithms Developed for Lodovico

Starting from the created datasets and others available in the literature, Deep Learning architectures for HTR have been developed. This research activity was based on the emerging Deformable Convolutional Network, which were originally designed for object recognition and segmentation (cf. Cojocaru et al., 2021 [32]). These types of algorithms are robust to non-idealities of paper support and are able to dynamically adapt to variations in the shape and size of characters. Before the appearance of our approach, DefConvs have indeed been employed for the task of object recognition, showing great adaptability to geometric variations and to part deformations, and the ability to model transformations in the object scale, pose, and viewpoint. To the best of our knowledge, we have been the first to investigate the usage of of DefConvs for handwriting recognition. Their kernel adaptability, indeed, helps to improve the efficiency and the performance in the task. We refer the reader to (Cojocaru et al., 2021 [32]) for further details on the architecture and training modality, but a qualitative example is shown in Figure 5, on test set lines of the IAM (Marti et al., 2022 [33]) and RIMES (Augustin et al., 2006 [34]) datasets, in comparison to the approach proposed by Shi et al. (Shi et al., 2016 [35]), a popular approach for HTR (see also Figure 6).

Moreover, in the context of the Lodovico project, we also developed a demo of the automatic annotation results. This is publicly available (https://ailb-web.ing.unimore.it/hwr, accessed on 20 November 2023), under the name of "Transcribe" platform (Figure 7). In addition to showcasing results, the platform can also be used to annotate new data. Finally, preliminary tests of text recognition from the pages of the Cronache Modenesi by Spaccini were also conducted using the developed models, and are available on the interface.

Parallel to the development of neural architectures specifically designed for HTR, work has been carried out concerning the training strategies aimed at improving their performance while maintaining or even reducing the amount of manually annotated data required for training. In this regard, the "Leopardi" dataset was chosen as a case study, and the focus was on the semi-automatic collection of synthetic and specific training data. Specifically, a randomized typeset font was created based on genuine glyphs commonly produced by Giacomo Leopardi to capture his graphic style. The font was used to transcribe the author's own prose works to capture his literary style and language. The transcribed text was overlaid onto an image of an empty page from the same collection of letters that make up the dataset to capture the characteristics of the type of paper support commonly used by the author to write the documents in the dataset. Finally, an automatic line-level annotation of the synthetic texts was obtained as described, automatically isolating the line images within the pages. This procedure enabled the creation of a synthetic dataset of over 100,000 samples, making it at least an order of magnitude larger than the HTR datasets available in the literature. The synthetic dataset was used for pre-training HTR models,

which were then directly applied to text recognition in the "Leopardi" dataset or fine-tuned on a small portion of real data. Particularly, in the latter case, the benefits of the proposed strategy of pre-training on carefully designed synthetic data were evident.

Figure 5. Some receptive fields of the developed deformable convolutional network (in red) and a standard convolutional network with the same architecture but without utilizing deformable convolutional operators (in blue). The deformable receptive fields are non-contiguous and irregularly shaped areas that adapt to pixels containing written strokes and cover a larger area of the image, providing more context while maintaining a compact number of parameters compared to the case of the standard convolutional network.
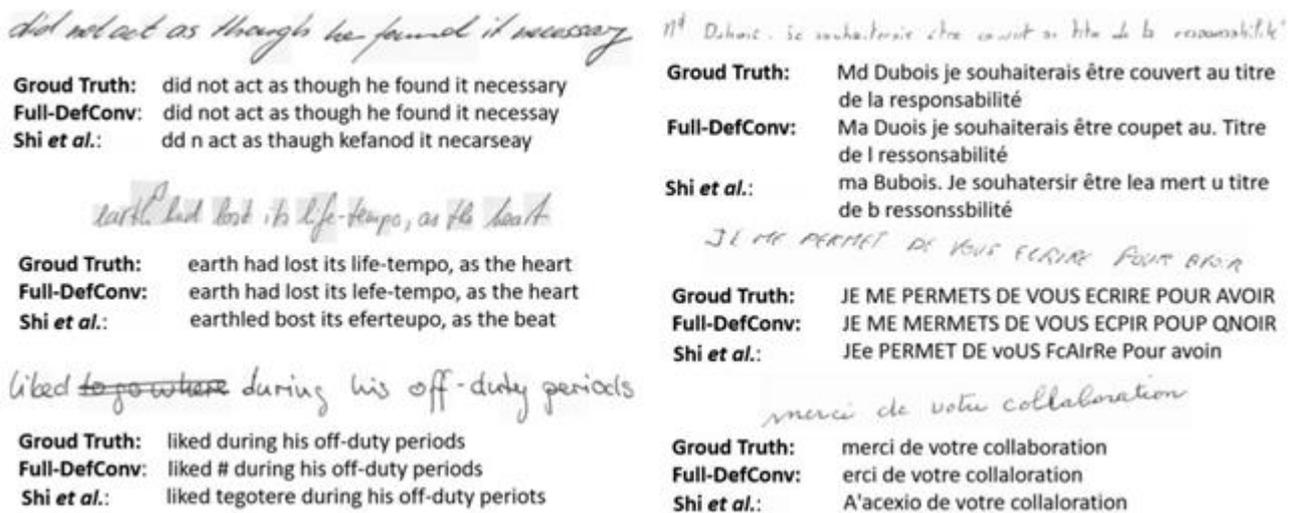

Figure 6. Qualitative handwritten recognition results on some test lines of the IAM dataset (**left**) and the RIMES dataset (**right**) [35].
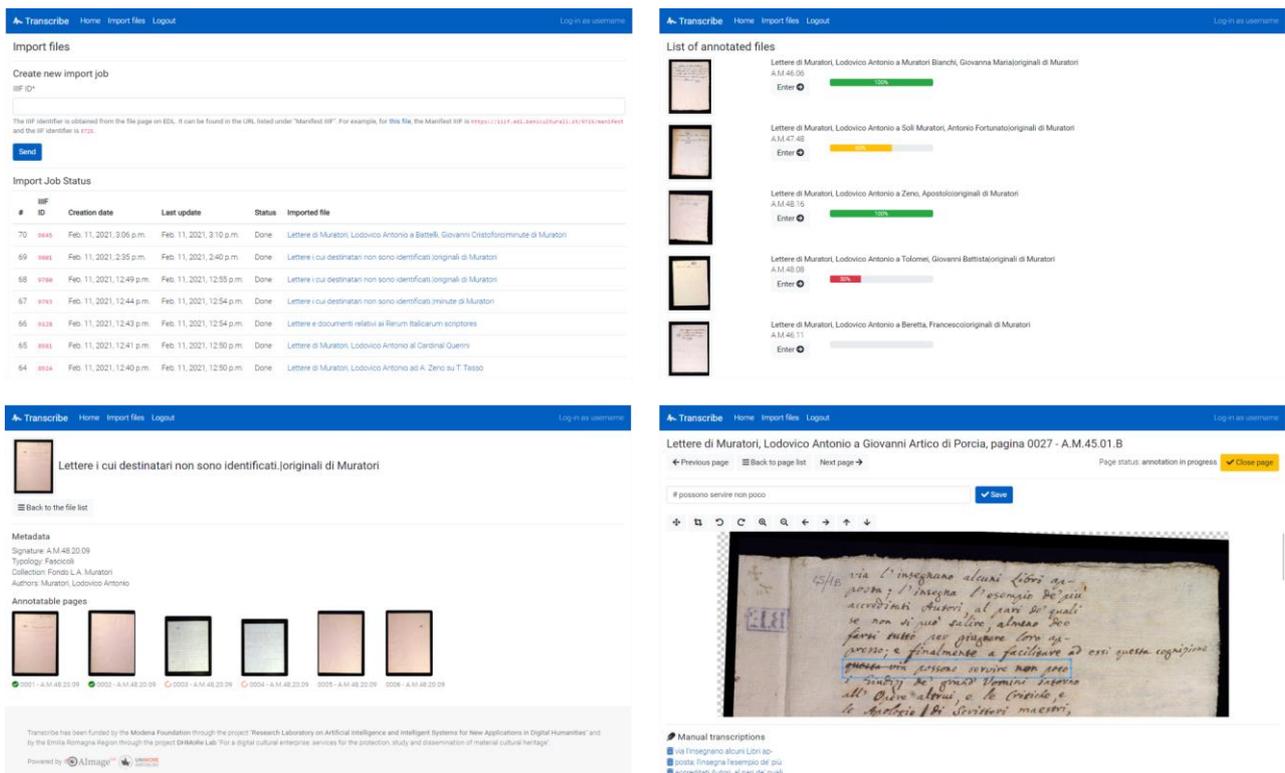
Figure 7. Screenshots of the "Transcribe" annotation platform developed for the collection of the "Leopardi" and "LAM" datasets.

Some quantitative results of this approach are reported in Table 2, where we also report the results of the direct application of the models pretrained on Leopardi Synth for comparison. The results are showcased in terms of their Character Error Rate or CER (which is proportional to the number of wrongly recognized characters) and Word Error Rate or WER (which instead considers the correctness of the transcription at the word level). The benefits of pretraining are more evident in the case of fine-tuning on 50% of the training lines. In fact, in this case, the CER decreases by 1.8 and the WER of 4.1 on average, while in the case of fine-tuning on 100% of the training lines, the CER decreases by 0.8 and the WER of 3.1 on average. Overall, these results demonstrate the effectiveness of training on synthetically generated data and open up new opportunities in terms of training HTR networks in low-data regimes.

Table 2. Experimental results of the considered models when pretrained on the Leopardi synthetic dataset and then fine-tuned on different portions of the real Leopardi dataset, compared to their performance when trained from scratch on the same lines of the real Leopardi dataset.

	Shi et al., 2016 [35]		Puigcerver et al., 2017 [36]		Shi et al., 2016 [35] + DefConv		Puigcerver et al., 2017 [36] + DefConv	
	CER	WER	CER	WER	CER	WER	CER	WER
ICFHR14	47.2	102.7	59.7	120.1	77.8	112.4	103.9	147.1
ICFHR16	76.0	129.8	79.1	111.0	83.1	109.3	86.9	144.6
IAM	46.5	92.9	68.0	96.4	60.4	97.5	74.4	98.9
RIMES	43.4	88.0	73.9	103.9	72.8	100.0	69.5	97.1
Leopardi	35.9	86.4	38.5	93.9	40.2	92.1	36.3	94.0

As the Lodovico library continues to expand, further explorations will involve the design of novel HTR algorithms capable of operating in scenarios with limited data and adapting to different handwriting styles. The ultimate goal is to integrate these algorithms

into the digital library, serving as a valuable tool for experts and the curious alike, facilitating enhanced navigation and search experiences.

5. Final Remarks

To conclude, it is possible to offer some final remarks concerning the results and potential of the Lodovico project. As we have seen, Lodovico does not aim to be an aggregator nor the digital library of a cultural institute: the challenge was to create a media library that integrates, in a single system, different cultural organisations belonging to a territory with homogeneous characteristics (in this case the Emilia-Romagna region).

Three main means were used to achieve this objective:

- a common and potentially interoperable data architecture (based on the Dublin Core standard);
- a metadata standardisation method shared by all partners;
- the overcoming of barriers concerning the various types of cultural heritage (towards a cross-typological media library).

Such an approach makes it easier to reconstruct the unity of historically dismembered and dispersed heritages and, more generally, favours the comparison of digital objects and the reconstruction of histories based on complex cultural heritage.

The federative logic thus generates new knowledge, not only thanks to the potential given by the increased accessibility of cultural heritage, but also through the combination ab origine of different heritages in a unified framework. In this regard, the first experiments carried out by Lodovico have already produced interesting results, revealing, for example, the hitherto unknown connections between documents stored in different cultural institutions.

Lodovico's potential lies, finally, in its collaborative premises: the project partnership is, in fact, coordinated by the university and, therefore, makes use of the expertise of the research world; at the same time, it is made up of institutes that preserve cultural heritage and, therefore, keeps the theoretical questions of research constantly anchored to the needs of conservation and dissemination of the heritage itself, in a communicative perspective. This double level—the world of research and cultural institutions—also makes it possible to reuse digitisations for experimental purposes, as demonstrated by the HTR software developed by the researchers of the Lodovico project.

Author Contributions: Conceptualization, M.A.K. and L.B.; methodology, M.A.K. and L.B.; software, L.B.; validation, M.A.K. and L.B.; formal analysis, M.A.K. and L.B.; investigation, M.A.K. and L.B.; resources, M.A.K. and L.B.; data curation, M.A.K. and L.B.; writing—original draft preparation, M.A.K. (Sections 2 and 3) and L.B. (Section 4); writing—review and editing, M.A.K. (Sections 2 and 3) and L.B. (Section 4); visualization, M.A.K. (Sections 2 and 3) and L.B. (Section 4); supervision, M.A.K. (Sections 2 and 3) and L.B. (Section 4); project administration, M.A.K. and L.B. All authors have read and agreed to the published version of the manuscript.

References

1. Burdick, A.; Lunenfeld, P.; Drucker, J. *Digital_Humanities*; MIT Press: Cambridge, MA, USA; London, UK, 2012.
2. Hotson, H.; Wallnig, T. *Reassembling the Republic of Letters in the Digital Age: Standards, Systems, Scholarship*; Göttingen University Press: Göttingen, Germany, 2019.
3. Edelstein, D.; Findlen, P.; Ceserani, G.; Winterer, C.; Coleman, N. Historical Research in a Digital Age: Reflections from the Mapping the Republic of Letters Project. *Am. Hist. Rev.* **2017**, *122*, 400–424. [CrossRef]
4. Imbruglia, G. Muratori, Ludovico Antonio. In *Dizionario Biografico Degli Italiani*; Istituto della Enciclopedia Italiana: Rome, Italy, 2012; p. 77.
5. Presner, T.; Johanson, C. The Promise of Digital Humanities: A Whitepaper. Available online: https://humanitiesblast.com/Promise%2520of%2520Digital%2520Humanities.pdf (accessed on 20 November 2023).

6. Levenberg, L.; Neilson, T.; Rheams, D. (Eds.) *Research Methods for the Digital Humanities*; Palgrave MacMillan: Cham, Switzerland, 2018.

7. Signorotto, G.; Tongiorgi, D. (Eds.) *Modena Estense. La Rappresentazione Della Sovranità*; Edizioni di Storia e Letteratura: Rome, Italy, 2018.

8. Fumagalli, E.; Signorotto, G. (Eds.) La Corte Estense Nel Primo Seicento. In *Diplomazia e Mecenatismo Artistico*; Viella: Rome, Italy, 2012.

9. Folin, M. Rinascimento Estense. In *Politica, Cultura, Istituzioni di un Antico Stato Italiano*; Laterza: Rome, Italy; Bari, Italy, 2001.

10. Marini, L. *Lo Stato Estense*; UTET: Torino, Italy, 1979.

11. Di Pietro Lombardi, P. (Ed.) *L'antico Fondo Della Biblioteca Estense Universitaria di Modena: I Manoscritti Latini (1–200)*; Istituto poligrafico e Zecca dello Stato-Libreria dello Stato: Rome, Italy, 2017.

12. Bonacini, P. Il Registrum Comunis Mutine, 1299. In *Politica e Amministrazione Corrente del Comune di Modena Alla Fine del XIII Secolo*; Archivio Storico: Modena, Italy, 2002.

13. Borsari, A. *La Memoria Della Città: l'Archivio Storico del Comune di Modena*; Archivio Storico, Comune, Assessorato Alla Cultura: Modena, Italy, 2001.

14. Biondi, A. Per una storia dell'attività consiliare nel comune di Modena dal medio evo alla fine dell'antico regime (1796). In *I Registri Delle Deliberazioni Consiliari del Comune di Modena dal XIV al XVIII Secolo*; Liotti, C., Romagnoli, P., Eds.; Coptip: Modena, Italy, 1987; pp. 7–43.

15. Bonacini, P. Terre d'Emilia. In *Distretti Pubblici, Comunità Locali e Poteri Signorili Nell'esperienza di Una Regione Italiana, Secoli VIII–XII*; Clueb: Bologna, Italy, 2001.

16. Vigarani, G.; Baldelli, F. (Eds.) *Inventario dei Manoscritti Dell'archivio Capitolare di Modena*; Poligrafico Mucchi: Modena, Italy, 2003.

17. Golinelli, P. *Nuova Storia Illustrata di Modena*; Pacini: Pisa, Italy, 2011.

18. Previtali, G. *Che Cosa Sono le Digital Humanities*; Carocci: Rome, Italy, 2022.

19. Ernst, W. *Digital Memory and the Archive*; University of Minnesota Press: Minneapolis, MN, USA, 2012.

20. Burgio, V. Rumore Visivo. In *Semiotica e Critica Dell'infografica*; Mimesis: Milan, Italy, 2021.

21. Spiro, L. This is why we fight: Defining the Value of the Digital Humanities. In *Debates in the Digital Humanities*; Gold, M.K., Ed.; University of Minnesota Press: Minneapolis, MN, USA, 2012; pp. 12–35.

22. Mahomi, S.; Pierazzo, E. Teaching Skills or Teaching Methodology? In *Digital Humanities Pedagogy: Practices, Principles and Politics*; Hirsch, B.D., Ed.; Open Book Publishers: Cambridge, UK, 2012; pp. 215–225.

23. Al Kalak, M.; Fumagalli, E. (Eds.) *Collezionare Autografi. La Raccolta di Giuseppe Campori*; Olschki: Florence, Italy, 2022.

24. Baraldi, L.; Cornia, M.; Grana, C.; Cucchiara, R. Aligning text and document illustrations: Towards visually explainable digital humanities. In Proceedings of the International Conference on Pattern Recognition, Beijing, China, 20–24 August 2018; pp. 1097–1102.

25. Cascianelli, S.; Cornia, M.; Baraldi, L.; Piazzi, M.L.; Schiuma, R.; Cucchiara, R. Learning to Read L'Infinito: Handwritten Text Recognition with Synthetic Training Data. In Proceedings of the International Conference on Computer Analysis of Images and Patterns, Virtual, 28–30 September 2021; pp. 340–350.

26. Toselli, A.; Juan, A.; Gonz, J.; Salvador, I.; Vidal, E.; Casacuberta, F.; Keysers, D.; Ney, H. Integrated handwriting recognition and interpretation using finite-state models. *Int. J. Pattern Recognit. Artif. Intell.* **2014**, *18*, 519–539. [CrossRef]

27. Graves, A.; Schmidhuber, J. Offline Handwriting Recognition with Multidimensional Recurrent Neural Networks. *Adv. Neural Inf. Process. Syst.* **2008**, *21*, 545–552.

28. Zhang, Y.; Nie, S.; Liu, W.; Xu, X.; Zhang, D.; Shen, H.T. Sequence-to-sequence domain adaptation network for robust text image recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 2740–2749.

29. Wigington, C.; Stewart, S.; Davis, B.; Barrett, B.; Price, B.; Cohen, S. Data augmentation for recognition of handwritten words and lines using a CNN-LSTM network. In Proceedings of the International Conference on Document Analysis and Recognition, Kyoto, Japan, 9–15 November 2017; Volume 1, pp. 639–645.

30. Zhong, Z.; Zhang, X.Y.; Yin, F.; Liu, C.L. Handwritten Chinese character recognition with spatial transformer and deep residual networks. In Proceedings of the International Conference on Pattern Recognition, Cancun, Mexico, 4–8 December 2016; pp. 3429–3434.

31. Bhunia, K.; Das, A.; Bhunia, A.K.; Kishore, P.S.R.; Roy, P.P. Handwriting Recognition in Low-resource Scripts using Adversarial Learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 4767–4776.

32. Cojocaru, I.; Cascianelli, S.; Baraldi, L.; Corsini, M.; Cucchiara, R. Watch Your Strokes: Improving Handwritten Text Recognition with Deformable Convolutions. In Proceedings of the International Conference on Pattern Recognition, Milan, Italy, 10–15 January 2021; pp. 6096–6103.

33. Marti, U.-V.; Bunke, H. The iam-database: An english sentence database for offline handwriting recognition. *Int. J. Doc. Anal. Recognit.* **2002**, *5*, 39–46. [CrossRef]

34. Augustin, E.; Carre, M.; Grosicki, E.; Brodin, J.-M.; Geoffrois, E.; Preteux, F. RIMES evaluation campaign for handwritten mail processing. In Proceedings of the International Workshop on Frontiers in Handwriting Recognition, La Baule, France, 23–26 October 2006; pp. 231–235.

35. Shi, B.; Bai, X.; Yao, C. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *IEEE Trans. PAMI* **2016**, *39*, 11. [CrossRef] [PubMed]
36. Puigcerver, J.; Toselli, A.H.; Vidal, E. Querying out-of-vocabulary words in lexicon-based keyword spotting. *Neural Comput. Appl.* **2017**, *28*, 2373–2382. [CrossRef]