

Towards a Digital Critical-Genetic Edition of the Autograph Fragments of *Orlando Furioso*: Methods, Tools, and Workflow

Ambra Torregrossa
Università eCampus

Abstract

This contribution presents a digital pipeline for the genetic edition of the autograph fragments of *Orlando Furioso*. Developed within a doctoral research project, this framework treats Ariosto's manuscript simultaneously as a philological, documentary, and computational object. The objective extends beyond the mere migration of established analogue content into electronic formats; rather, it aims to construct a research environment capable of preserving the material complexity of the folios, representing the stratification of textual variants, and ensuring that the editorial process remains verifiable at every stage. In this perspective, high-definition digitization, interoperable image management, the training of custom Handwritten Text Recognition (HTR) models, XML-TEI encoding, and open-source web publication converge into a unified methodological architecture. Within this system, technology serves as a tool for formalizing philological data rather than as a substitute for critical interpretation.

1 Introduction

Despite a highly distinguished editorial tradition, the scholarly reception of Ariostean manuscripts has long been constrained by the material limitations of print culture. This article takes as its point of departure Santorre Debenedetti's milestone 1937 edition, which remains an indispensable benchmark for the study of the *Orlando Furioso* manuscripts. Concurrently, however, Debenedetti's work exemplifies an editorial mediation inevitably conditioned by the static nature of the codex and by the attendant necessity to linearize writing processes that are, in reality, simultaneous and stratified.

The print format inherently privileges the chronological succession of corrections over their

topographical distribution across the page, thereby fragmenting the physical unity of the witness and relegating to separate columns or apparatus pages material that coexists within the same graphic space in the manuscript.

Consequently, the necessity for a digital edition stems from a strictly philological imperative: to restore the systemic relationship among the text, its material support, the writing gesture, and the compositional process, without collapsing these distinct layers into a one-dimensional representation. Within this framework, technology is not deployed as an automated panacea, but rather as a suite of procedures capable of expanding the editor's descriptive capacities – provided it is underpinned by rigorous, explicit, and verifiable criteria.

The very notion of the scholarly edition is thus operationally redefined: the reader must be enabled to reconstruct editorial choices, assess their empirical foundations, and, within the parameters allowed by the extant documentation, verify their soundness at both the material and interpretive levels.

On this basis, the project operates at the intersection of digital and computational philology. In the present approach, the former encompasses editorial practices conceived natively within the digital environment – characterized by multimedia integration, interactivity, hypertextuality, and layered data structures. The latter incorporates advanced computational methods, such as machine learning algorithms applied to computer vision and natural language processing, which play an increasingly critical role in integrating OCR/HTR workflows for text processing, modeling, and large-scale analysis.

The edition of Ariosto's autograph fragments is situated precisely at the nexus of these two domains, as it demands both the design of a native digital publication environment and the implementation of trainable models for the recognition of historical scripts.

2 The Corpus and the Problems of Representation

The pipeline operates on a composite manuscript corpus originating from a unified compositional practice. Its core component is the manuscript preserved at the Biblioteca Comunale Ariostea in Ferrara (shelfmark Cl. I A It. 1), a paper folio codex consisting of 54 leaves organized into ternary or quaternary quires. This core is supplemented by two autograph leaves from the Biblioteca Ambrosiana in Milan (formerly part of a single sheet and currently preserved in the *Cartella ariostesca*, S.P. 33), and one sheet plus an additional autograph leaf housed at the Biblioteca Nazionale Vittorio Emanuele III in Naples (ms. San Martino 353). These primary witnesses are accompanied, as auxiliary reconstructive material, by an apograph copy of the Ambrosiana manuscript (ms. H 55 inf.), which documents advanced stages of transcription and revision.

The significance of this corpus lies not only in its status as direct evidence of the author's hand, but in the opportunity it affords to trace Ariosto's compositional labor across a series of episodes destined for the final 1532 edition of the poem in 46 cantos. These include sections relating to Olympia, the Rocca di Tristano, Marganorre, Ruggiero and Leone, and *Lo scudo della regina Elisa*.

Furthermore, the documentation illuminates the variety of material supports utilized by the author at successive stages of composition: loose sheets and half-sheets were employed for fluid drafts and heavily revised textual materials, whereas structured quires were reserved for more advanced drafts, typically laid out according to the autograph norm of five *ottave* per page, leaving wide lateral margins for additions and corrective interventions.

This material configuration engenders at least three orders of editorial challenges:

- **Codicological Reconstruction:** Re-establishing the original sequence of the folios, particularly where internal lacunae, insertions, or material displacements noted by

scholarship complicate the macro-structure.

- **Genetic Representation:** Rendering the various stages of textual elaboration, ranging from rough drafts (*brouillons*) to corrected fair copies, up to text forms closely approaching the printed version.
- **Digital Architecture:** Ensuring the coexistence of highly heterogeneous data within a single, integrated system – including high-definition digital facsimiles, codicological descriptions, diplomatic transcriptions, semantic markup, genetic relationships, critical apparatuses, and navigation tools.

The pipeline was specifically engineered to address this textual complexity without resorting to reductive oversimplifications.

3 General Architecture of the Pipeline

The project's methodology relies on a source-to-output model that proceeds in a logical sequence: from the machine-readable transcription of the text, through semantic encoding according to the XML-TEI standard, to the design of the Graphical User Interface (GUI) dedicated to the publication of the digital edition. Rather than a linear assembly line, this sequence constitutes a circular validation system wherein each phase informs and constrains the others: image quality conditions the accuracy of the transcription; the transcription guides data modeling; the encoding schema restricts visualization choices; and the visualization interface, in turn, necessitates a redefinition of the semantic granularity of the markup.

The organizing principle governing this workflow is the primacy of the source file. The XML-TEI document acts as a centralized knowledge base and structured data repository from which multiple outputs can be dynamically generate – such as HTML for web rendering, PDF or ePub for portable reading, or varied configurations of the same text (e.g., diplomatic vs. critical-interpretive views) tailored to diverse scholarly requisites.

This single-source approach prevents data duplication and intra-textual inconsistencies, as all revisions are executed upstream within the master file and subsequently propagated downstream across the various publication formats. Such an approach

is paramount in philological scholarship, where the reliability of an edition depends entirely on the traceability of the transformations applied to the text.

This architecture further implies a dual dimension of control:

- **The Front-End:** Renders the text, critical apparatuses, facsimiles, and analytical tools accessible to the end-user.
- **The Back-End:** Exposes the underlying encoding logic, the editorial conventions adopted, and the schemas governing data representation.

Transparency is therefore not merely entrusted to a narrative statement of editorial criteria, but is embedded within the formal, machine-readable legibility of the system that implements them.

4 Imaging, Preservation and Interoperability

The foundational layer of the pipeline consists of the digital reproductions of the manuscripts. The custodial libraries play a strategic role in this workflow, as both the execution of high-definition digitization and the long-term archiving of the master files depend on their infrastructure.

The project adopts the FITS (Flexible Image Transport System) format as the standard solution for preserving master images. FITS ensures the integrity of complex, high-resolution scientific data and embeds rich metadata concerning image technical specs, object description, digitization parameters, provenance, and intellectual property rights within the file header, thereby satisfying strict criteria for traceability and long-term digital preservation.

For access and dissemination purposes, however, static master files are insufficient. Here, the IIIF (International Image Interoperability Framework) framework is deployed, utilizing a standardized set of APIs that enable seamless access, delivery, and interoperable visualization of digital images. Via IIIF-compliant servers, users can remotely examine high-definition facsimiles, manipulate zoom levels, compare folios from separate repositories side-by-side, and integrate them into compatible web viewers without the need to download large original files.

Consequently, FITS serves as the secure “digital vault” safeguarding the original high-resolution data, while IIIF acts as the “virtual museum” interface. This setup allows researchers to display faithful digital reproductions on screen, magnify them to scrutinize paleographical minutiae, and collate them seamlessly – even when the physical components are scattered across geographically distant institutions.

From the standpoint of authorial philology (*filologia d'autore*), the methodological benefits are profound: the verifiability of the editor's reading is guaranteed by the absolute synchronicity of the encoded text and the digital image, while the physical dispersion of the corpus is overcome through virtual reunification within a unified digital environment.

5 Training HTR Model's for Ariosto's Hand

Following image acquisition, the workflow transitions to the transcription phase and the generation of the ground truth utilizing the *Transkribus* platform. Developed under the aegis of the European READ project, *Transkribus* serves as a virtual research infrastructure (VRI) encompassing the entire data-production cycle, from layout segmentation to automated text recognition.

In processing the autograph fragments of *Orlando Furioso*, the primary methodological challenge lies not in the application of generic, pre-existing HTR models, but in training custom models specifically calibrated to Ludovico Ariosto's handwriting. The idiosyncratic graphic profile of the autograph – characterized by dense abbreviations, variations in *ductus*, interlinear corrections, marginal additions, and complex palimpsest-like rewriting – necessitates a highly precise alignment between the documentary data and the recognition algorithms.

To achieve this, a training sample consisting of 40 meticulously segmented and transcribed pages was compiled, from which *Transkribus* automatically extracted a validation set comprising 10% of the total corpus. Ground-truth preparation requires the rigorous definition of text regions, baselines, and baseline polygons, given that the fundamental learning unit for the PyLaia HTR engine is the statistical association between the graphic image of a text line and its corresponding validated transcription string.

In this context, layout segmentation is far from a mechanical preliminary step; it represents a critical variable determining the overall accuracy of the text recognition phase. The training workflow is structured around two complementary models:

1. **The Text Model:** Optimized for the recognition and transcription of verbal and alphanumeric content.
2. **The Field Model:** Tailored to identify, segment, and classify the distinct structural and topographical components of the page layout.

This bifurcated approach is especially vital for genetic criticism, as it separates the paleographical problem of reading the text from the structural problem of mapping its spatial distribution on the support, thereby allowing the system to capture linguistic and topographical data simultaneously.

The evaluation metrics provided by the platform – specifically the convergence curves of the training and validation loss for the field model, and the Character Error Rate (CER) for the text model – constitute essential diagnostics for assessing the models' learning trajectory, generalization capacity, and output quality. These quantitative indicators do not replace human philological oversight; rather, they enable the editor to diagnose systematic errors, refine the training sample, and establish a statistically grounded confidence level for the automatically recognized text. In this regard, both trained models achieved excellent evaluation metrics following several fine-tuning experiments performed on the public pre-trained model for Italian (e.g., *Transkribus Italian Handwriting MI*)

Machine learning thus functions as an assistive technology subject to continuous human validation, fully integrated into a methodology centered on data verifiability and the editor's critical responsibility.

6 XML-TEI as a Semantic Model

The core phase of the pipeline entails encoding the corpus according to the Text Encoding Initiative (XML-TEI) guidelines, the established international standard in Digital Humanities for the semantic representation of textual sources. Diverging from models oriented toward the mere graphic or typographic replication of a document, TEI describes textual phenomena via a hierarchical, tree-

structured architecture that renders their function, significance, and mutual relationships explicit, transforming the text into an interoperable, queryable database resource.

Furthermore, adherence to this standard guarantees the formal validation of data through rigorous schemas, enforces a shared domain vocabulary, and ensures the long-term digital sustainability of the project's assets by relying on an open, non-proprietary format.

While a comprehensive analysis of the complete encoding schema is beyond the scope of this paper, it is crucial to outline how the architecture has been designed to support a native digital genetic edition of a poetic work, synthesising bibliographic, codicological, paleographical, metrical, and authorial-genetic data layers.

The primary TEI modules utilized and their respective functions are as follows:

- **teiHeader:** Captures descriptive, administrative, and bibliographic metadata of the digital edition.
- **msdescription:** Documents the material aspects of the witness (support, structure, page layout, state of preservation, and provenance history).
- **namesdates:** Facilitates the prosopographical and geographical normalization of historical actors, locations, and chronological expressions.
- **transcr:** Forms the backbone of the *representation of primary sources*, recording additions, deletions, substitutions, changes of ink/hand, and topographical text distribution.
- **textcrit:** Manages textual variance and the critical-genetic apparatus, formalizing successive redactional stages and structural relationships between variant readings.
- **verse:** Preserves and annotates the formal metrical and stanzaic architecture of the poem (*ottava rima*).
- **facsimile:** Establishes direct, coordinate-based links mapping transcription

segments to specific zones on the manuscript images.

- **gaiji**: Encodes and documents non-standard glyphs, scribal allographs, and historical ligatures.

The encoding schema is applied directly to a diplomatic transcription that rigorously preserves historical orthography, punctuation, capitalization, abbreviations, and other material features of the graphemic system. This produces an integrated data model capable of reconciling the material concreteness of the source with the fluid dynamics of the author's compositional process.

7 From Encoding to Visualization

Once the data has been structured and validated, the focus shifts to publication infrastructure, revealing a methodological divergence between two potential web-delivery solutions: *Transkribus Sites* and EVT.

The first option, *Transkribus Sites*, offers a user-friendly, turn-key interface natively integrated into the *Transkribus* ecosystem. Ideal for medium-scale initiatives, it facilitates the online deployment of images, aligned text-image transcriptions, metadata, full-text search engines, and basic document navigation without requiring dedicated, custom web development.

However, this platform presents non-trivial limitations for advanced scholarly workflows: the user interface design offers restricted customization, the infrastructure relies on a commercial, third-party service provider, and the editor retains only partial sovereignty over the underlying source code and data management. These factors carry direct, adverse implications for long-term editorial autonomy and scientific reproducibility.

Conversely, the preferred infrastructure for this project is EVT (Edition Visualization Technology), an open-source, highly adaptable visualization environment developed at the University of Pisa under the direction of Roberto Rosselli Del Turco. Built entirely on standard web technologies, fully compliant with modern web browsers, and designed around a modular framework, EVT ensures a higher degree of transparency, sustainability, and institutional independence than its commercial alternatives.

The trade-off associated with this second path involves increased technical overhead, as the

deployment of EVT demands advanced digital competencies for local installation, software configuration, resource indexing, and server-side optimization. For a highly specialized philological enterprise, however, this granular level of architectural control is vital: it allows the digital interface to be shaped by the idiosyncratic requirements of the textual corpus, rather than forcing the corpus to conform to the constraints of a pre-fabricated system.

8 Current Results and Analytical Limitations

The outcomes achieved to date are significant, particularly regarding the successful training of custom AI models optimized for Ariosto's complex layout and handwriting, alongside the definition of a robust, scalable prototype encoding model. Nonetheless, these results remain partial due to the inherent complexity of the project and its deeply interdisciplinary nature, which requires the continuous convergence of high-level legal, philological, computational, and project-management expertise.

On a practical level, this initiative has successfully achieved the digital reintegration of the missing autograph leaves preserved at the Biblioteca Ambrosiana, empirically demonstrating how a digital ecosystem can virtually reconstruct the unity of a physically fragmented archival dossier.

Despite this success, substantial gaps persist in the digital reconstruction due to the current unavailability of certain key collateral materials. These include the apograph copy relating to the Rocca di Tristano, the textual fragment *Per la storia d'Italia*, and three crucial autograph leaves from the Biblioteca Nazionale di Napoli pertaining to the Ruggiero and Leone episodes and the *Lo scudo della regina Elisa*.

This shortfall leads to an important methodological conclusion: a philologically sound digital pipeline depends as much on institutional coordination as it does on the sophistication of its computational tools. Inter-institutional cooperation among research libraries constitutes an absolute prerequisite for ensuring open access to primary materials, standardizing high-quality reproductions, and guaranteeing the long-term sustainability of the digital edition. Absent such institutional synergy, even the most advanced technical environment remains structurally incomplete, as technology can neither substitute for the physical availability of primary sources nor synthetically

compensate for structural lacunae in the documentary record.

9 Concluding Considerations

The research experience detailed in this paper demonstrates that when digital philology confronts an autograph dossier characterized by high genetic density, it cannot be reduced to passive digitization or to the merely decorative deployment of computing tools. Its scientific specificity lies in its capacity to transform the editorial process into a formalized chain of documented, reversible, and machine-verifiable operations, wherein the digital facsimile, the text, the markup, and the user interface collectively converge to construct a single, unified critical object.

In the case of the Ariostean fragments, the pipeline engineered around FITS, IIF, *Transkribus*, XML-TEI, and EVT demonstrates that digital media can simultaneously capture the graphic materiality of the manuscript, the shifting dynamics of authorial revision, and the structured, semantic legibility of the underlying data. The scientific advantage consists not merely in broader public access, but in the capacity to represent – without reductive compression – the simultaneous coexistence of material, textual, and genetic layers that traditional print culture fails to preserve synoptically.

This model underpins a rigorous definition of the digital scholarly edition: it is not a mere screen-based publication of a text, but a dynamic knowledge system wherein every critical choice is explicitly exposed, every tier of the text is fully searchable, and every visualization variant references back to a semantically structured source database. In this sense, the Ariostean project offers a scalable methodological framework adaptable to other authorial corpora, particularly those where material dispersion, complex scribal stratification, and the need to integrate philological rigor with computational expertise render a unified, accurate, and sustainable pipeline indispensable.

References

- Albonico, Simone. 2019. “Pour un retour aux manuscrits de l’Arioste.” *Génésis* 49: 31–46. Reprinted as “Dalla parte di Ariosto” in *Gli ‘scartafacci’ degli scrittori. I sentieri della creazione letteraria in Italia*, edited by Christian Del Vento and Paolo Musi-telli, 143–62. Rome: Carocci, 2022.
- Albonico, Simone. 2022. “Ariosto.” In *Autografi dei letterati italiani. Il Cinquecento*, edited by Matteo
- Motolese, Paolo Procaccioli, and Emilio Russo, 3:3–35. Rome: Salerno Editrice.
- Ariosto, Ludovico. 1937. *I frammenti autografi dell’Orlando furioso*. Edited by Santorre Debenedetti. Turin: Chiantore. Reprinted with a preface by Cesare Segre. Rome: Edizioni di Storia e Letteratura, 2010.
- Burnard, Lou. 2014. *What Is the Text Encoding Initiative?* Marseille: OpenEdition Press.
- Campeggiani, Irene. 2017. *L’ultimo Ariosto. Dalle Satire ai Frammenti autografi*. Pisa: Scuola Normale Superiore.
- Mancinelli, Tiziana, and Elena Pierazzo. 2020. *Che cos’è un’edizione scientifica digitale*. Rome: Carocci.
- Mühlberger, Günter, et al. 2019. “Transforming Scholarship in the Archives through Handwritten Text Recognition: Transkribus as a Case Study.” *Journal of Documentation* 75 (5): 954–76.
- Pierazzo, Elena. 2014. “Digital Documentary Editions and the Others.” *Scholarly Editing: The Annual of the Association for Documentary Editing* 35.
- REED-COOP. 2024. “Transkribus Sites – Publishing Documents Made Easy.” *Transkribus Blog*, January 10, 2024.
- Rosselli Del Turco, Roberto. 2019. “Designing an Advanced Software Tool for Digital Scholarly Editions: The Inception and Development of EVT (Edition Visualization Technology).” *Textual Cultures* 12 (2): 91–111.
- Sahle, Patrick. 2016. “What Is a Scholarly Digital Edition?” In *Digital Scholarly Editing: Theory and Practice*, edited by Matthew James Driscoll and Elena Pierazzo, 19–40. Cambridge: Open Book Publishers.
- TEI Consortium. 2025. *TEI P5: Guidelines for Electronic Text Encoding and Interchange* (P5 Version 4.9.0, last updated January 24, 2025).