

Pre-Editorial Normalization for Automatically Transcribed Medieval Manuscripts in Old French and Latin

Thibault Clérice,¹ Rachel Bawden,¹ Anthony Glaise,²
Ariane Pinche,³ David A. Smith⁴

¹Inria, Paris, France ²CNRS (UMR 7323), Tours, France,

³CNRS (UMR 5648), Lyon, France, ⁴Northeastern University

{firstname.lastname}@{inria,cnrs}.fr, dasmith@ccs.neu.edu

Abstract


Recent advances in Automatic Text Recognition (ATR) have improved access to historical archives, yet a methodological divide persists between palaeographic transcriptions and normalized digital editions. While ATR models trained on more palaeographically-oriented datasets such as CATMuS have shown greater generalizability, their raw outputs remain poorly compatible with most readers and downstream NLP tools, thus creating a usability gap. On the other hand, ATR models trained to produce normalized outputs have been shown to struggle to adapt to new domains and tend to over-normalize and hallucinate. We introduce the task of Pre-Editorial Normalization (PEN), which consists in normalizing graphemic ATR output according to editorial conventions, which has the advantage of keeping an intermediate step with palaeographic fidelity while providing a normalized version for practical usability. We present a new dataset derived from the CoMMA corpus and aligned with digitized Old French and Latin editions using *passim*. We also produce a manually corrected gold-standard evaluation set. We benchmark this resource using ByT5-based sequence-to-sequence models on normalization and pre-annotation tasks. Our contributions include the formal definition of PEN, a 4.66M-sample silver training corpus, a 1.8k-sample gold evaluation set, and a normalization model achieving a 6.7% CER, substantially outperforming previous models for this task.

Keywords: Latin, Old French, Normalization, Text Recognition, Medieval manuscripts

1. Introduction

The progress of Automatic Text Recognition (ATR) has facilitated access to historical archives (Pinche and Stokes, 2024). However, there exists a methodological divide in how to transcribe manuscripts on a spectrum ranging between two extremes: the graphetic approach and a normalized approach (Figure 1). The graphetic approach, defined by Robinson and Solopova (1993), treats transcription as a graphical representation of the manuscript’s text, which captures each character variant, and maintains original punctuation and abbreviation without expansion. Normalized transcription conflates transcribing with textual editing (Pierazzo, 2011), thus adhering more to users expectations by “[...] [making] the manuscript text accessible to the modern reader by expanding abbreviations, separating words, inserting diacritical marks, and adding modern punctuation” (Duval, 2012, p.330). However, this approach, combining two tasks at the same time, a graphical interpretation and a linguistic interpretation, has shown to poorly resolve abbreviations and to hallucinate content (Aguilar, 2024; Torres Aguilar and Jolivet, 2023), without providing an ability for the reader to distinguish normalization errors from graphical interpretation errors.

It can be useful to follow an approach between the two extremes, whereby character variants are merged but where abbreviations and other typi-



Graphetic	ĩmoðatis uenĩs ufĩbz nat̃alis calor
Graphemic	ĩmoðatis uenĩs usib; nat̃alis calor
Normalized	inmoderatis veneris usibus naturalis calor

Figure 1: The outputs of different transcription approaches for the same line of text. Manuscript: BSB CLM 13027. The task defined in this paper consists in producing normalized transcriptions from graphemic outputs.

cal manuscripts phenomena are preserved: this is known as the graphemic approach. This enables the transcriptions to be treated as their own epistemic objects, which is an advantage in itself. It also results in transcriptions that result in models, such as those trained on the CATMuS dataset (Clérice et al., 2024), that better generalize to new domains and languages (Coll Ardanuy et al., 2025).

However, graphemic transcriptions create a “usability gap” where the most accurate transcriptions of the material reality of the text are the least compatible with most available tools for computational analysis and expectations from the humanities (Kiessling, 2025).¹ It can therefore be use-

¹Tools such as lemmatizers, part-of-speech (PoS) taggers, and Named Entity Recognition (NER) systems

ful to automatically produce normalized versions of these transcriptions. It holds an advantage over all-in-one normalized ATR approaches, as it provides an intermediate layer that makes it possible to distinguish ATR error (such as <rn>/<m>) from normalization errors (<ē&&> for <est et> OR <esset>).

In this paper, we introduce the task of automatically producing normalized transcriptions from graphemic ATR transcriptions, which we name pre-editorial normalization (PEN). The task definition and rationale is followed by the introduction of a novel dataset, based on the ATR silver corpus CoMMA (Clérice et al., 2025) and digitized and born-digital editions of Old French and Latin texts. Using *passim* (Smith et al., 2015), we automatically aligned both corpora, of which we then manually corrected some of the alignments to produce a gold-standard test set. We benchmark the silver dataset using models trained on a sequence-to-sequence transformer, ByT5 (Xue et al., 2022), across two tasks: (i) a normalization task, for which we consider our gold test set as the goal, and (ii) a pre-annotation task for lemmatization and PoS-tagging.

2. Related Work

2.1. Automatic Text Recognition with Normalized Output

ATR research has generally paid limited attention to the distinction between imitative transcription of historical sources and their subsequent normalization. Historical text recognition models are often evaluated on resources such as the IAMHistDB medieval datasets, the Perzival (Fischer et al., 2009), and the Saint Gall Database (Fischer et al., 2011), despite the fact that these corpora rely on different transcription paradigms. Camps et al. (2022) benchmarked both approaches on genre-constrained documents but without providing detailed qualitative analyses of the results.

By contrast, several studies within the Digital Humanities community have explicitly addressed the risks associated with over-normalization. In particular, two recent works have examined this issue as a central research question. Bottaioli et al. (2024) show that models trained to jointly transcribe and normalize middle names in digitized birth certificates may hallucinate names for individuals who did not possess one. Similarly, in the context of medieval manuscripts and, more broadly, in languages with highly variable orthography, Tarterolo-Orta et al. (2025) demonstrate that increased ac-

are typically trained on normalized, modern-style editions. When faced with the raw, non-resolved output of a graphemic ATR model, the performance of these tools drops significantly (Clérice and Pinche, 2025).

cessibility through normalization entails a latent risk of systematically erasing historically and linguistically meaningful variation.

Conversely, qualitative error analyses conducted in several studies by S. T. Aguilar highlight persistent limitations of models trained on normalized outputs, particularly with respect to abbreviation resolution. Torres Aguilar and Jolivet (2023) show that morphosyntactic errors constitute one of the most frequent sources of abbreviation-related errors. Moreover, Aguilar (2024) report several critical failures, including difficulties in handling line-end hyphenation, the misinterpretation of simple abbreviations (e.g., <īfusū> for <infusum> transcribed as <in futurum>) even in relatively straightforward contexts, and errors in local semantic disambiguation, such as interpreting <dia³> as <diaconus> ‘deacon’ instead of <diabolus> ‘devil’.

2.2. Historical Text Normalization

Datasets for historical text normalization remain scarce, particularly for Latin. To our knowledge, the only comprehensive and readily available non-synthetic dataset for Latin historical normalization is the work of Schonhardt (2025a,b), which is limited to a single text attested in six manuscripts. This restricted scope limits its suitability for studying abbreviation expansion in a broader diachronic and bilingual context. A small number of additional datasets have been produced or identified in the work of Camps et al. (2022), some of which include Old French. However, their coverage remains limited in terms of genre and textual diversity.

Related work by Büttner (2025) addresses the normalization of Latin spelling by aligning medieval forms with more classical orthographic standards using a ByT5-based approach. While this study provides a relevant example of normalization in the Latin domain, it focuses on orthographic regularization rather than pre-editorial normalization.

By contrast, normalization for vernacular languages has received greater attention, beginning with late Middle French and extending into the early modern period. Solfrini et al. (2025) present a study on the normalization of printed late Middle French texts from the 1530s. More broadly, the automatic standardization of late Middle French has been explored in several works (Rubino et al., 2024b,a). These approaches are similar in methodology but differ in purpose from ours, as their primary aim is to reduce endemic linguistic variation rather than to preserve philological and historical features. On this topic, Bawden et al. (2022) compare normalization models for early modern French (17th c. and later), including datasets and comparative assessments of statistical and neural machine translation approaches.

Finally, digital editions may encode such information using a variety of representation frameworks, as illustrated in doctoral work (Franzini, 2019) and other projects (Nury, 2024). However, beyond the technical challenges involved in interpreting and exporting data from these platforms, identifying suitable editions remains a difficult task. To date, the only available catalogue providing filters for downloadable TEI resources (Andorfer and Franzini, 2025) indicates that, among more than 29 digital editions with accessible XML-TEI files for Latin and French texts, none preserves abbreviated forms with explicit markup.²

2.3. Normalization of User-Generated Content (UGC)

Another type of data associated with the lexical normalization task is user-generated content (UGC), of which social media texts are the most prevalent (Baldwin et al., 2013). UGC, like the historical data dealt with in this article, is characterised by non-standard phenomena such as acronyms, abbreviations, truncations and variations in spelling. Such texts can present challenges for natural language processing (NLP) tools, particularly those trained on standard, edited texts, which has led to a body of work in normalisation as a preprocessing step (Melerio et al., 2012; Baldwin et al., 2015; van der Goot et al., 2021; Nishimwe et al., 2024a).

Initial approaches relied on rules, lexicons and language models, and it was common to decompose the task into the detection of words to be modified and their normalisation (Baranes, 2015; Ikeda et al., 2016). Later approaches opt for a translation-like approach, where the task is treated as the translation of a non-standard to standard language sequence (Aw et al., 2006; De Clercq, Orphée and Desmet, Bart and Schulz, Sarah and Lefever, Els and Hoste, Veronique, 2013). This approach has the advantage of increased use of context with the sequence and flexibility, helpful notably for segmentation and reordering differences.

The choice of representation has been a subject of much research in the handling of UGC texts (both for normalisation and other NLP tasks). Phonetic representations have been explored (Rosales Núñez et al., 2019; Mandal and Nanmaran, 2018), particularly for spelling errors. Character representations have also shown promise in terms of increasing robustness against phenomena such as spelling errors, devowelling, acronyms and abbreviations (Belinkov and Bisk, 2018; Watson et al., 2018; Domingo and Casacuberta, 2021; Scherrer and Ljubešić, 2021). Character-level representations have the advantage of better capturing similarities between similar but non-identical lexical

²Sahle et al. (2020) does not provide such metadata.

forms. It is an approach that has previously been shown to work well for morphologically rich languages (which are also characterised by a higher level of variation) and low-resource scenarios (in order to generalize better across forms). For this reason, another successful approach is to either collect and annotate more data or to generate synthetic data (Belinkov and Bisk, 2018; Vaibhav et al., 2019; Nishimwe et al., 2024b).

Recent work has shown that models trained on large quantities of data, and in particular large language models (LLMs) are already “more robust” in that they perform better than previously on non-standard UGC data, without being specifically trained to do so and without having to apply a specific normalization pre-processing step (Bawden and Sagot, 2023; Peters and Martins, 2025; Bawden and Sagot, 2025).³ However, this is unlikely to be the case across all languages, and in particular low-resource ones.

2.4. Automatic Alignment

Aligning multiple data representations has been an important method for dataset creation since the application of dynamic time warping in speech recognition (Sakoe and Chiba, 1978) or dynamic programming for aligning ATR outputs or multiple translations of the same source (Pial and Skiena, 2023; Thai et al., 2022). This paper uses methods developed for text-reuse detection, which is a broad task with many possible approaches and evaluations (Moritz et al., 2016). As described below (§4.2), the main inference task in creating our dataset is to match and align the ATR output of manuscript pages with passages in normalized editions. Unlike in other ATR evaluation tasks, these normalized editions do not have page breaks that match any particular manuscripts and the manuscripts are often miscellanies with several works. The size of the corpus makes it impractical to compare all pairs of ATR pages and passages from digital editions (Paju et al., 2022). We therefore employ passim 2.0 (Smith et al., 2023), which uses hashed character n -gram fingerprinting to filter the set of candidate passages for alignment. Once a pair of passages is selected, passim uses a character-level hidden Markov alignment model (Vogel et al., 1996). Unlike text-reuse packages that use the Smith-Waterman alignment algorithm (Smith et al., 2014; Funk and Mullen, 2018), the HMM approach is more robust to the rearrangement of passages in the ATR output due to uncertainty in page layout analysis.

³Given that they are trained on large amounts of web data, it can be debated whether the models are actually more robust or whether the non-standard forms are simply sufficiently common in the training data.

3. Pre-editorial Normalization

3.1. Task Definition

The output of graphemic ATR⁴ is inherently raw and reflects both manuscript-specific phenomena and transcription guidelines. As a result, it exhibits several sources of variability: spacing is inconsistent (including but not limited to unmarked line-break hyphenation), capitalization—especially in earlier centuries—may serve a decorative rather than linguistic function, numerical expressions vary across systems (ranging from Roman numerals to fractional notation in later periods), ramist letter pairs (<u/v> and <i/j>) are not systematically distinguished, punctuation practices frequently differ from those adopted in modern editorial standards, and abbreviations are preserved. Finally, especially in the context of raw ATR output, recognition errors are superimposed on top of the multiple sources of variation described above.

These characteristics considerably limit the direct usability of ATR outputs in standard NLP and text-processing pipelines. Moreover, they diverge substantially from the normalized representations typically expected by scholars in the humanities, thereby motivating the need for dedicated normalization and post-processing methods.

The PEN task consists in:

- normalizing letters that were considered allo-graphs in the middle ages and that diverged over time, such as ramist letters (<u>/<v>) in consonantal and vocalic usage. Historical spelling variants are not normalized; for instance, <f> replacing an etymological <ph> is preserved.
- resolving abbreviations in a context-sensitive manner: expansions should reflect not only the semantic content of words, but also their morphosyntactic features and a coherent orthographic system within the text.
- addressing recognition errors in the context of raw ATR output, such as <rn>/<m> confusions.
- normalizing capitalization around punctuation and proper names according to modern editorial standards.
- preserving original punctuation.

⁴Robinson and Solopova (1993) defines graphemic as a transcription keeping original practices such as abbreviation but merging variants of same characters, such as <s> and <ʃ>. They are implemented and more broadly defined in Clérice et al. (2024).

3.2. Limits of the Task

Certain abbreviations, especially those referring to personal names or place names, should not be expanded when the available knowledge in context is insufficient to support a reliable interpretation. For example, the abbreviation <G. dixit> ‘G. said’ in a medical text can be reliably expanded to <Galenus dixit> based on contextual evidence. On the other hand, in a charter or accounts, there might not be any contextual elements to resolve the name of S. in <S. Parisiensis> ‘S. the Parisian’, which could be many things (Stephanus, Severinus, etc.) and the initial should therefore be kept as it is.

Severe or unprocessable ATR errors resulting in extended spans of noise should not be corrected by the model. Attempting to normalize such segments risks introducing over-normalization and, more critically, hallucinated content. Given the text-based nature of the task, only recoverable textual information should be considered.

Numerical expressions written in Roman or Arabic numerals are likewise left unchanged. This form of variation remains an open issue within the community and requires further standardization before being fully integrated into the task definition. Moreover, the distinction between ordinal and cardinal numbers is often expressed through contextual or morphological markers and may require complex morphosyntactic analysis to be resolved reliably on top of the inherent variation of Roman numeral spelling.

Sufficient contextual information is required at all levels (semantic, morphosyntactic, and spelling) to support reliable normalization decisions. However, contextual evidence may itself be conflicting. For instance, if both <molt> and <moult> occur within the same Old French text, it remains unclear how the abbreviated form <mlt> should be expanded.

In this paper, we frame the task as a single input–single output problem. In future work, the PEN task should also involve a processing layer capable of representing a distribution over possible outputs, to account for spelling variations within a single document, which can be seen in Old French (<segnor> and <signor> on the same page).

4. Dataset Construction

4.1. Source Data

Our data for historical transcriptions is drawn from the CoMMA corpus, a graphemic-transcription corpus of 33,000 manuscripts (3.3B whitespace-delimited tokens). In our approach, we use the main text of each page to build full-text manuscript pages, ignoring running headers, marginal zones, and interlinear comments, thanks to the dataset’s

fine-grained details concerning layout classification. Library manuscript records carry language metadata; however, as this metadata operates at the document level, it is not used for filtering — our alignment process may identify passages in Latin or Old French within manuscripts bearing no such annotation.

For the normalized texts to align with CoMMA, we use publicly available data for Latin, and a mix of closed and open data for Old French, to ensure a large enough pool of data to align with the manuscripts. Both the Latin pool and the Old French one are drawn from the same data cited by CoMMA for their ModernBert model (Clérice et al., 2025, Table 6). Some of the corpus contains highly technical data (mathematics, medical, etc.), which we expect to be the hardest to normalize in Latin, as it contains higher levels of special vocabulary and specific abbreviations.

4.2. Alignment

We align the ATR transcripts of the CoMMA corpus with the plain text of digital editions of Latin and Old French texts using passim 2.0 (Smith et al., 2023). This version of the package implements a character-level HMM alignment algorithm (Vogel et al., 1996), which permits rearrangement of passages between two versions of a text. This is helpful in working with ATR output, when the sequence of page regions in the transcript does not match the sequence as transcribed in a scholarly edition. In contrast to other ATR ground-truth datasets produced by forced alignment, we do not need to specify in the input which manuscripts are versions of which texts; rather, we specify to passim that each page from CoMMA should be aligned to one or more pages from the corpus of digital editions. We performed unicode decomposition normalization on both the CoMMA ATR output and the digital editions so that the common use of diacritics for, e.g., nasalization (⟨cōsul⟩ for ⟨consul⟩) or truncation (⟨c̄⟩ for ⟨cum⟩) would align more closely to their standard forms. To improve recall, we index the collection with character 10-grams (ignoring whitespace and punctuation) that occur in fewer than 100 pages. It then aligns pages having at least five 10-grams in common and drops alignments shorter than 50 characters. We use beam search ($b = 200$) when aligning pages.

From the automatic alignment, we filter the data with at least five lines of continuous aligned content, a sufficiently large match rate (60%), and only where the alignment comes from the same source text. This resulted in 3.2M alignments. We then reprocessed them into smaller alignments, with a maximum and minimum size of 1000 and 300 bytes respectively, to get down to a total of 4,665,509 pairs.

Corpus	Work	# Mss.
PL	Auctor Incertus - Missale mixtum	2090
PL	Gratianus - Concordia discordantium canonum	2026
encyMe	Speculum doctrinale	2002
CSEL	Augustine - Liber qui appellatur Speculum	1795
PL	Martinus Legionensis - Sermones	1735
PL	Auctor incertus - Breviarium	1625
PL	Ivo Carnotensis - Decretum	1613

Table 1: Top works by manuscript presence.

4.3. Limits

As our approach relies on the automatic alignment of raw ATR output with digitized editions, several sources of noise may affect the resulting annotations regarding the task definition. First, since the distance between an abbreviated token and its expanded form influences alignment quality, there is a non-negligible probability that a manuscript form is aligned with a different lexical variant in the critical edition. This situation is common in textual transmission and may result in semantically or syntactically divergent alignments. Additionally, over-normalization may also occur, in a manner similar to variant-swapping phenomena, such as ⟨f⟩/⟨ph⟩ alternations in Latin or ⟨moult⟩/⟨molt⟩ variation in Old French. These cases may lead to normalized forms that contradict the intended task definition.

4.4. Description

Unsurprisingly, classical, late-antique, and Christian “literary” Latin texts are over-represented in the alignments between our source corpora and CoMMA, while later medieval administrative texts are under-represented. This reflects the fact that CoMMA is largely derived from library collections rather than archives. Patrologia Latina, our largest source (97M tokens, Roelli and Ctibor (2024)), accounts for 64.5% of the alignment pairs. By contrast, CEMA, an archival corpus and the second-largest source (57M tokens, Perreaux (2021)), represents only 1.5% of the alignments, whereas Perseus (6.1M tokens, Crane et al. (2021b)) accounts for 6.5%.

As expected from studies on manuscript transmission, some works are much more represented than others (see Table 1), and all most common authors are either tied to religious practice (missals, sermons and homilies), theory (commentaries), or canon law (*decretum*). One reason for this over-representation of some texts might be the presence of numerous bible citations.

Over-normalization, in the form of variant insertion, spelling standardization (in Latin), or lexical replacement (in Old French), constitutes one of the main identified risks. At this stage, however, we are unable to reliably distinguish over-




Figure 2: Distribution of the percentage of tokens per alignment in the train set that do not correspond to abbreviation resolution pattern.

normalization from two related phenomena: first, cases in which abbreviation markers are not recognized by the ATR system (e.g. <ifusa> for <īfusā> with missed diacritics), and second, post-ATR editorial corrections, which are likely more prevalent.

Nevertheless, we estimate the extent of over-normalization by automatically aligning tokens in the training set and measuring the proportion of tokens that differ between the source and aligned texts when no explicit abbreviation markers are present. This analysis reveals a marked contrast between the two languages (Figure 2). In Latin, 20.5% of tokens exhibit substitutions consistent with non-standard abbreviation resolution. This proportion broadly aligns with expectations based on the 9.7% character error rate (CER) reported for CoMMA, including additional effects from the limited possibilities of over-normalization in Latin and variant substitutions.

By contrast, Old French displays a substantially higher level of non-abbreviation related substitutions, with a median substitution rate of 51% (and a mean of 50%). This discrepancy suggests that, compared to Latin, Old French editions impose a greater degree of spelling normalization on manuscript sources, thereby introducing more extensive spelling variation between manuscripts and their corresponding editorial representations.

4.5. Manual Correction and Gold Standard

To create a gold set, we sampled some data to be representative of main corpora from the automatic alignment. The resulting pairs are all manually corrected by a philologist to resolve alignment errors, normalization inconsistencies, and segmentation issues. This results in a gold-standard dataset

		Word Length			
		Samples	Average	Std	Median
Raw ATR	French	189	95	34	89
	Latin	352	93	39	88
GT ATR	Latin	45	288	256	259
	Mixed	2	276	101	276

Table 2: Distribution of languages, sample counts, and token counts across the two alignment sources: raw ATR output and ground truth (GT).

that serves both for evaluation and for qualitative error analysis. To that test data, we add ground truth data from CATMuS data repositories (both archives and literary documents) that was manually normalised by the same expert.

In the automatically aligned ATR output-part of the gold dataset, two texts are substantially over-represented: *Oresme* (75 alignments) and the *Histoire Ancienne* (43). This imbalance reflects the limited diversity of available Old French datasets, which made it challenging to construct a fully balanced sample. All other corpora are represented by between 11 and 21 alignments. Among these, some technically demanding texts are included, such as the *Mathematica* corpus. It is represented by 20 alignments and contains atypical elements such as geometrical point labels (e.g. ‘A’), which resemble abbreviations but must not be treated as such.

The ground truth ATR focused specifically on two series of texts: the first are drawn from archives, small letters and charters, mostly from the Cartularies of Notre-Dame de Paris (Claustre and Smith, 2022); the second are drawn from a single manuscript bearing the text *Liber de Coitu* (Constantine the African, x1 c.). The first series contains administrative text, with various anonymous people and occasional switching between Latin and French, while the second contains very abbreviated content, such as “eliq̄t̄ hūor q̄ ē.ī cēb° ȝ eliq̄t̄?” (corresponding to *eliquatur humor qui est in cerebro, et eliquatus* ‘the humor that is in the brain liquefies, and once it is liquefied’).

The final gold dataset comprises 588 samples (64,230 tokens) from 316 different manuscripts. 47 alignments were drawn from ground truth ATR, with no or very few recognition errors, and which are therefore longer than other samples (287 tokens on average vs. 94, see Table 2).

5. Benchmarking Normalization

Our first experiment evaluates the impact of the proposed dataset in the context of ATR output normalization. Specifically, we train a model for abbreviation resolution and compare its performance

against off-the-shelf approaches, including zero-shot and domain-limited pretrained models.

5.1. Setup

We formulate text normalization as a sequence-to-sequence task. We select ByT5 as the base architecture, as it can handle a wide range of characters (including Unicode-decomposed forms), since it operates at the byte level. This property has been associated with improved performance on languages exhibiting high levels of spelling variation (Wolters and Van Cranenburgh, 2024). ByT5 has also been successfully applied to related tasks such as ATR correction (Pavlopoulos et al., 2024), medieval Latin normalization (Büttner, 2025), and abbreviation expansion (Schonhardt, 2025a).

Baselines We consider two baseline systems. The first is a ByT5-based model trained on the dataset of Schonhardt (2025a). This model relies on different Unicode-level guidelines for abbreviations and is strongly domain-constrained. Despite the shared scribal conventions for abbreviation between Old French and Latin, we expect this model to perform poorly on Old French, which is out of domain. In addition, the dataset’s source text were manually corrected, which may hinder robustness to ATR noise and to inconsistencies in spacing.

Our second baseline consists of zero-shot prompting Llama 3 8B Instruct (Grattafiori et al., 2024). The model is prompted to perform PEN, with an auxiliary instruction to predict the input language, and common abbreviation patterns and glyphs (Appendix D).

Training data The training data is bilingual, combining Latin and Old French. Owing to a strong imbalance between the two languages and the relative scarcity of Old French data, we upsample the French portion of the training set by a factor of ten, reducing the imbalance from approximately 100:1 to 10:1 (Latin to Old French), and thereby ensuring sufficient exposure to Old French during training.

Experimental design We evaluate on our manually curated test set, as well as on the first 2,000 entries of the dataset introduced by Schonhardt (2025a), for which the Schonhardt baseline model should be at an advantage. All experiments are evaluated using CER and word error rate (WER), using the Levenshtein edit distance normalized over the length of the ground truth.

5.2. Results

Evaluated on our gold dataset, our proposed model performs well for both languages, improving over the zero-shot model and the only other

Model	CER			WER		
	All	Latin	Old French	All	Latin	Old French
<i>On our gold data</i>						
Llama 3	139.9	160.4	79.9	163.7	183.1	107.3
ByT5 (Schonhardt)	64.8	75.6	33.3	92.3	104.8	55.9
ByT5 (Ours)	6.7	6.3	7.6	20.7	19.6	23.8
<i>On Schonhardt data (Latin only)</i>						
Llama 3		26.4			54.5	
ByT5 (Schonhardt)		130.0			109.4	
ByT5 (Ours)		11.1			38.5	

Table 3: Comparison of PEN results.

available model for the task by significant margins (Table 3). On the Schonhardt dataset, our model shows a stronger performance than the original model itself, the original model showing a tendency to undergenerate.

Qualitative analysis reveals a recurrent *classical bias* in the normalization of Latin, whereby the model tends to favor forms aligned with standardized spelling over historically attested medieval variants, except for final <-e> instead of <ae>. This tendency is illustrated by systematic preferences for <ph> over <f> and for <-tio> over <-cio>, reflecting a reversion to classical spellings. It also overpredicts names without context knowledge, such as <Johannis> for <Io.> and often struggles to adapt abbreviation resolution to the context, not unlike some of the models of Torres’ work.

In Old French, the absence of stabilized orthographic conventions similarly leads the model to privilege dominant variants in its training set, normalizing forms such as <leurs> to <lors> and <pueple> to <pople>. This results in the systematic erasure of localized *scripta* traditions, as the model prioritizes alignment targets over historically meaningful scribal variation. It however shows a strong capacity to retokenize or correct wrong ATR predictions.

Finally, we observe cases of variant insertion by the model, or at least semantic failure such as <iustum> for <iterum>. While this phenomenon constitutes a limitation of the baseline model introduced in this study, it nevertheless highlights the importance of decoupling PEN from the ATR process itself. From a data-centric perspective, this separation enables us to trace normalization errors back to pre-editorial stages without requiring direct recourse to the original manuscript sources.

6. Impact on a Downstream Task

Given these results and the observed overnormalization effects, a remaining question is whether our normalization model—trained on the alignment dataset—can nevertheless outperform a previously established pipeline for downstream text analysis. To address this question, we compare our approach with the pipeline proposed by Clérice and Pinche (2025) in the context of stylometric au-

Pipeline Name	Lemma			POS			POS 3-grams		
	Pre.	Rec.	F1	Pre.	Rec.	F1	Pre.	Rec.	F1
<i>All tokens</i>									
Lambert	88.2	89.4	88.8	95.8	97.2	96.5	85.6	86.8	86.2
Ours	90.9	89.0	89.9	98.0	95.9	96.9	87.7	85.8	86.7
<i>100 most frequent words (MFW)</i>									
Lambert	92.0	92.6	92.3						
Ours	94.3	89.5	91.8						

Table 4: Impact of the normalizer on downstream tasks (PoS tagging and lemmatization).

thorship analysis of *Vie de Saint Lambert*, focusing on the extraction of lemmas and PoS information, including PoS 3-grams.

6.1. Setup

Stylometric analysis of Old French texts is strongly hindered by spelling variation, particularly in inter-manuscript settings. Their pipeline addresses this issue by combining a convolutional neural network for space error correction with a model for both PoS-tagging and lemmatization (Clérice et al., 2019) fine-tuned on artificially abbreviated data (hereafter, the *Lambert* pipeline). On the other hand, our pipeline uses our normalization model and the same PoS-tagger/lemmatizer from (Clérice et al., 2019).

Data We conduct our evaluation on their transcription of *Vie de Saint Lambert* in the manuscript BnF Fr. 412. We compare their lemmatization and PoS-tagging outputs and the annotations produced by our own pipeline against the ground truth annotation of the normalized version. The *Vie* displays a very low amount of contextual abbreviation, and is mostly hindered by spacing issues and ATR errors. We compare the ground-truth based on the full results and on the 100 most frequent words (MFW), in a setting similar to stylometry practices (ignoring any punctuation).

Evaluation To assess the impact of these different preprocessing pipelines, we adopt a Bag-of-Words (BoW) evaluation framework and compute precision and recall in order to measure their suitability for frequency-based analyses. Given G and P the sequences of gold and predicted labels, and $c_G(x)$ and $c_P(x)$ the number of times a label x occurs in G and P , we define the number of true positives as $TP = \sum_x \min(c_G(x), c_P(x))$. We then have $\text{Precision}_{\text{BoW}} = \frac{TP}{|P|}$, and $\text{Recall}_{\text{BoW}} = \frac{TP}{|G|}$.

6.2. Results

Overall, both models produce a comparable number of tokens. The Lambert pipeline generates approximately 2% more tokens, whereas the ByT5-

based approach results in 3.8% fewer tokens than the edited text. For both systems, our error analysis reveals that annotation errors are not directly correlated with the correctness of the normalization step applied to individual tokens. For instance, the word “uie” (life), which is a non-ambiguous input for manual lemmatization, is poorly lemmatized 66% of the time by the Lambert pipeline and 30% by ours, with two out of three errors due to a non-normalization to “vie” and a lemmatization to $u\text{IE}$ instead of $v\text{IE}1$.

The results (Table 4) exhibit a consistent general trend across all evaluated subsets. Precision and F1 are systematically higher for our approach, within a few percentage points, while recall is consistently higher for the Lambert pipeline by a comparable margin (except for MFW lemmatization). Our model introduces less noise than Lambert’s pipeline but undergenerates and produces inadequate content for the tagger: for example, it does not produce apostrophes (“qu il” instead of “qu’il”), whereas they seem to be an overused marker by the tagger.

In the context of stylometric analysis, both approaches therefore present complementary advantages, depending on their respective impact on the distribution of the 100 most frequent words. The advantage of our approach lies in a potential simplification of the pipeline, contrary to the multiplication of steps and ad-hoc models from Lambert’s.

7. Discussion and Conclusion

We introduced pre-editorial normalization (PEN) as an intermediate task between graphemic ATR output and fully edited (normalized) text, and presented a large-scale aligned dataset for medieval Latin and Old French. Our experiments show that models trained on this resource substantially outperform existing baselines on normalization, and that PEN shows promising results as a first step for downstream linguistic analysis while simplifying processing pipelines.

At the same time, our results highlight persistent risks of over-normalization and editorial bias, particularly in the tendency to favor more classical spellings for Latin and alternative ones for Old French. These findings underline the importance of treating normalization as an intermediate layer rather than an implicit part of transcription.

Future work will focus on expanding the dataset, modeling uncertainty in normalization, and more importantly page-wide and document-level normalization settings, in line with recent practices in machine translation to use broader contextual windows. We hope that this resource and task definition will support more transparent and reusable approaches to research on historical texts.

8. Bibliographical References

- Sergio Torres Aguilar. 2024. [Handwritten text recognition for historical documents using visual language models and gans](#).
- Peter Andorfer and Greta Franzini. 2025. Catalogue digital editions. <https://dig-ed-cat.acdh.oeaw.ac.at/>.
- AiTi Aw, Min Zhang, Juan Xiao, and Jian Su. 2006. [A phrase-based statistical model for SMS text normalization](#). In *Proceedings of the COLING/ACL 2006 Main Conference Poster Sessions*, pages 33–40, Sydney, Australia. Association for Computational Linguistics.
- Timothy Baldwin, Paul Cook, Marco Lui, Andrew MacKinlay, and Li Wang. 2013. [How Noisy Social Media Text, How Diffrent Social Media Sources?](#) In *Proceedings of the 6th International Joint Conference on Natural Language Processing*, pages 356–364, Nagoya, Japan.
- Timothy Baldwin, Marie Catherine de Marneffe, Bo Han, Young-Bum Kim, Alan Ritter, and Wei Xu. 2015. [Shared tasks of the 2015 workshop on noisy user-generated text: Twitter lexical normalization and named entity recognition](#). In *Proceedings of the Workshop on Noisy User-generated Text*, pages 126–135, Beijing, China. Association for Computational Linguistics.
- Marion Baranes. 2015. [Normalisation orthographique de corpus bruités](#). Theses, Université Paris-Diderot - Paris VII.
- Rachel Bawden, Jonathan Poinhos, Eleni Kogkitsidou, Philippe Gambette, Benoît Sagot, and Simon Gabay. 2022. [Automatic normalisation of early Modern French](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3354–3366, Marseille, France. European Language Resources Association.
- Rachel Bawden and Benoît Sagot. 2023. [RoCS-MT: Robustness challenge set for machine translation](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 198–216, Singapore. Association for Computational Linguistics.
- Rachel Bawden and Benoît Sagot. 2025. [RoCS-MT v2 at WMT 2025: Robust challenge set for machine translation](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 834–849, Suzhou, China. Association for Computational Linguistics.
- Yonatan Belinkov and Yonatan Bisk. 2018. [Synthetic and natural noise both break neural machine translation](#). In *International Conference on Learning Representations*.
- Natalia Bottaioli, Solène Tarride, Jérémy Anger, Seginus Mowlavi, Marina Gardella, Antoine Tadros, Gabriele Facciolo, Rafael Grompone von Gioi, Christopher Kermorvant, Jean-Michel Morel, et al. 2024. [Normalized vs diplomatic annotation: A case study of automatic information extraction from handwritten uruguayan birth certificates](#). In *International Conference on Document Analysis and Recognition*, pages 40–54. Springer.
- Pascale Bourgain. 1992. [Sur l'édition des textes littéraires latins médiévaux](#). *Bibliothèque de l'École des chartes*, 150(1):5–49.
- Andreas Büttner. 2025. [byt5-base-latin-normalize: A latin orthography normalization model](#). <https://huggingface.co/andbue/byt5-base-latin-normalize>.
- Jean-Baptiste Camps, Chahan Vidal-Gorène, Dominique Stutzmann, Marguerite Vernet, and Ariane Pinche. 2022. [Data Diversity in handwritten text recognition. Challenge or opportunity?](#) In *Digital Humanities 2022. Responding to Asian diversity. Conference Abstracts (The University of Tokyo, Japan, 25-29 July 2022)*, Digital Humanities 2022. Responding to Asian diversity. Conference Abstracts (The University of Tokyo, Japan, 25-29 July 2022), pages 160–165, Tokyo, Japan. DH2022 Local Organizing Committee.
- Julie Claustre and Darwin Smith. 2022. [e-NdP notre-dame de paris et son cloître \(1326-1504\)](#). *Revue Mabillon*, 33:218–226.
- Thibault Clérice, Simon Gabay, Malamatenia Vlachou-Efstathiou, Ariane Pinche, and Benoît Sagot. 2025. [CoMMA, a Large-scale Corpus of Multilingual Medieval Archives](#).
- Thibault Clérice and Ariane Pinche. 2025. [Wauchier, Is That You? A multi-manuscript authorship analysis of Saint Lambert's life](#). In *CHR 2025 - 6th Conference on Computational Humanities Research*, Luxembourg, Luxembourg.
- Thibault Clérice, Ariane Pinche, Malamatenia Vlachou-Efstathiou, Alix Chagué, Jean-Baptiste Camps, Matthias Gille Levenson, Olivier Brisville-Fertin, Federico Boschetti, Franz Fischer, Michael Gervers, Agnès Boutreux, Avery Manton, Simon Gabay, Patricia O'connor, Wouter Haverals, Mike Kestemont, Caroline

- Vandyck, and Benjamin Kiessling. 2024. [CAT-MuS Medieval: A multilingual large-scale cross-century dataset in Latin script for handwritten text recognition and beyond](#). In *Lecture Notes in Computer Science*, volume 14806 of *Lecture Notes in Computer Science*, pages 174–194, Athens, Greece. Springer Nature Switzerland.
- Thibault Clérice, Jean-Baptiste Camps, and Ariane Pinche. 2019. [Deucalion, modèle ancien français \(0.2.0\)](#).
- Mariona Coll Ardanuy, Iban Berganzo-Besga, Ramon Sarobe, and Coral Cuadrada. 2025. Evaluating handwritten text recognition in medieval notarial manuscripts: A new dataset and comprehensive analysis. In *International Conference on Document Analysis and Recognition*, pages 340–357. Springer.
- De Clercq, Orphée and Desmet, Bart and Schulz, Sarah and Lefever, Els and Hoste, Veronique. 2013. Normalization of Dutch user-generated content. In *Proceedings of Recent Advances in Natural Language Processing*, pages 179–188. INCOMA.
- Miguel Domingo and Francisco Casacuberta. 2021. [A comparison of character-based neural machine translations techniques applied to spelling normalization](#). In *Pattern Recognition. ICPR International Workshops and Challenges: Virtual Event, January 10-15, 2021, Proceedings, Part VII*, page 326–338, Berlin, Heidelberg. Springer-Verlag.
- Frédéric Duval. 2012. [Transcrire le français médiéval: De l'« Instruction » de Paul Meyer à la description linguistique contemporaine](#). *Bibliothèque de l'École des chartes*, 170(2):321–342.
- Andreas Fischer, Volkmar Frinken, Alicia Fornés, and Horst Bunke. 2011. Transcription alignment of latin manuscripts using hidden markov models. In *Proceedings of the 2011 Workshop on Historical Document Imaging and Processing*, pages 29–36.
- Andreas Fischer, Markus Wuthrich, Marcus Liwicki, Volkmar Frinken, Horst Bunke, Gabriel Viehhauser, and Michael Stolz. 2009. Automatic transcription of handwritten medieval documents. In *2009 15th International Conference on Virtual Systems and Multimedia*, pages 137–142. IEEE.
- Greta Hayley Franzini. 2019. *Towards a first digital edition of the oldest surviving manuscript of St Augustine's De civitate Dei*. Ph.D. thesis, UCL (University College London).
- Kellen Funk and Lincoln A. Mullen. 2018. [The Spine of American Law: Digital Text Analysis and U.S. Legal Practice](#). *The American Historical Review*, 123(1):132–164.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Papsuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kam-badur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan,

Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Rapparthi, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gouget, Virginie Do, Vish Vogeti, Vitor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuwei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers,

Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabza, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Greshnev, Maxim Naumov, Maya Lathi, Meghan Kenally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Ritner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang,

- Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. 2024. [The Llama 3 Herd of Models](#).
- Bo Han and Timothy Baldwin. 2011. [Lexical normalisation of short text messages: Makn sens a #twitter](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 368–378, Portland, Oregon, USA. Association for Computational Linguistics.
- Taishi Ikeda, Hiroyuki Shindo, and Yuji Matsumoto. 2016. [Japanese text normalization with encoder-decoder model](#). In *Proceedings of the 2nd Workshop on Noisy User-generated Text (WNUT)*, pages 129–137, Osaka, Japan. The COLING 2016 Organizing Committee.
- Benjamin Kiessling. 2022. [The Kraken OCR system](#).
- Benjamin Kiessling. 2025. [Transcription Guidelines for Generalized Automatic Text Recognition](#).
- Minghao Li, Tengchao Lv, Jingye Chen, Lei Cui, Yijuan Lu, Dinei Florencio, Cha Zhang, Zhoujun Li, and Furu Wei. 2023. Trocr: Transformer-based optical character recognition with pre-trained models. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 13094–13102.
- Soumil Mandal and Karthick Nanmaran. 2018. [Normalization of transliterated words in code-mixed data using Seq2Seq model & Levenshtein distance](#). In *Proceedings of the 2018 EMNLP Workshop W-NUT: The 4th Workshop on Noisy User-generated Text*, pages 49–53, Brussels, Belgium. Association for Computational Linguistics.
- Héctor Martínez Alonso, Djamé Seddah, and Benoît Sagot. 2016. [From noisy questions to Minecraft texts: Annotation challenges in extreme syntax scenario](#). In *Proceedings of the 2nd Workshop on Noisy User-generated Text (WNUT)*, pages 13–23, Osaka, Japan. The COLING 2016 Organizing Committee.
- Claudia Matos Veliz, Orphee De Clercq, and Veronique Hoste. 2019. [Comparing MT approaches for text normalization](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, pages 740–749, Varna, Bulgaria. INCOMA Ltd.
- Maite Melero, Marta R. Costa-Jussà, Judith Domingo, Montse Marquina, and Martí Quixal. 2012. [Holaaa!! writin like u talk is kewl but kinda hard 4 NLP](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 3794–3800, Istanbul, Turkey. European Language Resources Association (ELRA).
- Paul Michel and Graham Neubig. 2018. [MTNT: A testbed for machine translation of noisy text](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 543–553, Brussels, Belgium. Association for Computational Linguistics.
- Maria Moritz, Andreas Wiederhold, Barbara Pavlek, Yuri Bizzoni, and Marco Büchler. 2016. [Non-literal text reuse in historical texts: An approach to identify reuse transformations and its application to Bible reuse](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1849–1859, Austin, Texas. Association for Computational Linguistics.
- Lydia Nishimwe, Benoît Sagot, and Rachel Bawden. 2024a. [Étude sur la normalisation lexicale de contenus produits par les utilisateurs \[a study on the lexical normalisation of user-generated content\]](#). *Traitement Automatique des Langues*, 64(2):15–41.
- Lydia Nishimwe, Benoît Sagot, and Rachel Bawden. 2024b. [Making sentence embeddings robust to user-generated content](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 10984–10998, Torino, Italia. ELRA and ICCL.

- Elisa Nury. 2024. [Codex bobbiensis \(vl 1\) dataset 8 \(f.36r to f.39v\)](#).
- Petri Paju, Heli Rantala, and Hannu Salmi. 2022. Towards an ontology and epistemology of text reuse. In Estelle Bunout, Maud Ehrmann, and Frédéric Clavert, editors, *Digitised Newspapers – A New Eldorado for Historians?: Reflections on Tools, Methods and Epistemology*. De Gruyter.
- Jungsoo Park, Mujeen Sung, Jinhyuk Lee, and Jaewoo Kang. 2020. [Adversarial subword regularization for robust neural machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1945–1953, Online. Association for Computational Linguistics.
- John Pavlopoulos, Vasiliki Kougia, Esteban Garces Arias, Paraskevi Platanou, Stepan Shabalin, Konstantina Liagkou, Emmanouil Papadatos, Holger Essler, Jean-Baptiste Camps, and Franz Fischer. 2024. [Challenging Error Correction in Recognised Byzantine Greek](#). In *Proceedings of the 1st Workshop on Machine Learning for Ancient Languages (ML4AL 2024)*, pages 1–12, Hybrid in Bangkok, Thailand and online. Association for Computational Linguistics.
- Ben Peters and Andre Martins. 2025. [Did Translation Models Get More Robust Without Anyone Even Noticing?](#) In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2445–2458, Vienna, Austria. Association for Computational Linguistics.
- Tanzir Pial and Steven Skiena. 2023. [GNAT: A General Narrative Alignment Tool](#). ArXiv:2311.03627 [cs].
- Elena Pierazzo. 2011. [A rationale of digital documentary editions](#). *Literary and Linguistic Computing*, 26(4):463–477.
- Ariane Pinche. 2021. *Edition nativement numérique du recueil hagiographique “Li Saint Confessor” de Wauchier de Denain d’après le manuscrit fr. 412 de la Bibliothèque nationale de France*. Ph.D. thesis.
- Ariane Pinche and Peter Stokes. 2024. [Historical documents and automatic text recognition: Introduction](#). *Journal of Data Mining & Digital Humanities*, Historical Documents and automatic text recognition.
- Peter Robinson and Elizabeth Solopova. 1993. Guidelines for transcription of the manuscripts of the wife of bath’s prologue. *The Canterbury Tales Project Occasional Papers*, 1:19–52.
- José Carlos Rosales Núñez, Djamé Seddah, and Guillaume Wisniewski. 2019. [Phonetic normalization for machine translation of user generated content](#). In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 407–416, Hong Kong, China. Association for Computational Linguistics.
- Raphael Rubino, Sandra Coram-Mekkey, Johanna Gerlach, Jonathan David Mutal, and Pierrette Bouillon. 2024a. [Automatic normalisation of Middle French and its impact on productivity](#). In *Proceedings of the Third Workshop on Language Technologies for Historical and Ancient Languages (LT4HALA) @ LREC-COLING-2024*, pages 176–189, Torino, Italia. ELRA and ICCL.
- Raphael Rubino, Johanna Gerlach, Jonathan Mutal, and Pierrette Bouillon. 2024b. [Normalizing without modernizing: Keeping historical word-forms of Middle French while reducing spelling variants](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 3394–3402, Mexico City, Mexico. Association for Computational Linguistics.
- Patrick Sahle, Georg Vogeler, Jana Klinger, Stephan Makowski, and Nadine Sutor. 2020. Digital scholarly editions. Online resource.
- Hiroaki Sakoe and Seibi Chiba. 1978. [Dynamic programming algorithm optimization for spoken word recognition](#). *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26:159–165.
- Yves Scherrer and Nikola Ljubešić. 2021. [Sesame street to mount sinai: BERT-constrained character-level Moses models for multilingual lexical normalization](#). In *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2021)*, pages 465–472, Online. Association for Computational Linguistics.
- Michael Schonhardt. 2025a. [Bdd abbreviations augmented: A parallel corpus for abbreviation expansion in medieval latin](#). [<https://huggingface.co/datasets/mschonhardt/bdd-abbreviations-augmented>] (<https://huggingface.co/datasets/mschonhardt/bdd-abbreviations-augmented>).
- Michael Schonhardt. 2025b. Die digitale edition als schnittstelle: Editorische workflows im mittelalter. *Das Mittelalter. Perspektiven mediävistischer Forschung*, 30(1):54–69.

- Djamé Seddah, Benoit Sagot, Marie Candito, Virginie Moulleron, and Vanessa Combet. 2012. [The French Social Media Bank: a treebank of noisy user generated content](#). In *Proceedings of COLING 2012*, pages 2441–2458, Mumbai, India. The COLING 2012 Organizing Committee.
- David A. Smith, Ryan Cordell, Elizabeth Maddock Dillon, Nick Stramp, and John Wilkerson. 2014. Detecting and modeling local text reuse. In *ACM-IEEE Joint Conference on Digital Libraries (JCDL)*.
- David A. Smith, Ryan Cordell, and Abby Mullen. 2015. [Computational Methods for Uncovering Reprinted Texts in Antebellum Newspapers](#). *American Literary History*, 27(3):E1–E15.
- David A. Smith, Jacob Murel, Jonathan Parkes Allen, and Matthew Thomas Miller. 2023. Automatic collation for diversifying corpora: Commonly copied texts as distant supervision for handwritten text recognition. In *Computational Humanities Research*.
- Sonia Solfrini, Mylène Dejoux, Aurélia Marques Oliveira, and Pierre-Olivier Beaulnes. 2025. [Normaliser le moyen français : du graphématique au semi-diplomatique](#). In *Actes des 18e Rencontres Jeunes Chercheurs en RI (RJCRI) et 27ème Rencontre des Étudiants Chercheurs en Informatique pour le Traitement Automatique des Langues (RECITAL)*, pages 239–252, Marseille, France. ATALA & ARIA.
- Katherine Thai, Marzena Karpinska, Kalpesh Krishna, Bill Ray, Moira Inghilleri, John Wieting, and Mohit Iyer. 2022. [Exploring Document-Level Literary Machine Translation with Parallel Paragraphs from World Literature](#). ArXiv:2210.14250 [cs].
- Sergio Torres Aguilar and Vincent Jolivet. 2023. [Handwritten Text Recognition for Documentary Medieval Manuscripts](#). *Journal of Data Mining and Digital Humanities*, Historical Documents and automatic text recognition.
- Yanco Amor Torterolo-Orta, Jaione Macicior-Mitxelena, Marina Miguez-Lamanuzzi, and Ana García-Serrano. 2025. [Transcribing Spanish Texts from the Past: Experiments with Transkribus, Tesseract and Granite](#).
- Vaibhav, Sumeet Singh, Craig Stewart, and Graham Neubig. 2019. [Improving robustness of machine translation with synthetic noise](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1916–1920, Minneapolis, Minnesota. Association for Computational Linguistics.
- Rob van der Goot, Alan Ramponi, Arkaitz Zubiaga, Barbara Plank, Benjamin Muller, Iñaki San Vicente Roncal, Nikola Ljubešić, Özlem Çetinoğlu, Rahmad Mahendra, Talha Çolakoğlu, Timothy Baldwin, Tommaso Caselli, and Wladimir Sidorenko. 2021. [MultiLexNorm: A shared task on multilingual lexical normalization](#). In *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2021)*, pages 493–509, Online. Association for Computational Linguistics.
- Stephan Vogel, Hermann Ney, and Christoph Tillmann. 1996. [HMM-based word alignment in statistical translation](#). In *Proceedings of the 16th Conference on Computational Linguistics (COLING)*, volume 2, page 836.
- Daniel Watson, Nasser Zalmout, and Nizar Habash. 2018. [Utilizing character and word embeddings for text normalization with sequence-to-sequence models](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 837–843, Brussels, Belgium. Association for Computational Linguistics.
- Andre Wolters and Andreas Van Cranenburgh. 2024. Historical dutch spelling normalization with pretrained language models. *Computational Linguistics in the Netherlands Journal*, 13:147–171.
- Linting Xue, Aditya Barua, Noah Constant, Rami Al-Rfou, Sharan Narang, Mihir Kale, Adam Roberts, and Colin Raffel. 2022. [ByT5: Towards a token-free future with pre-trained byte-to-byte models](#). *Transactions of the Association for Computational Linguistics*, 10:291–306.

9. Language Resource References

- Assmann, B. and Sahle, Patrick. 1826–. *Monumenta Germaniae Historica*.
- Bauduin, Pierre and Buard, Pierre-Yves and Goloubkoff, Anne and Mancel, Émeline and Fournier-Fujimoto, Tamiko and Buquet, Thierry and Brossard, Stéphane and Trohel, Benoît and Combalbert, Grégory and Bisson, Marie and Berthelot, Clémentine and Dubois, Adrien and Bloche, Michaël. 2023. *Scripta*.
- Jean-Baptiste Camps and Alice Cochet and Lucence Ing and Elena Albarran. 2019. [Geste: un corpus de chansons de geste, 2016-....](#) Zenodo.

- Clérice, Thibault. 2021. *Corpora of rare texts, Las-civa Roma*.
- Crane, Gregory Ralph and Babeu, Alison and Cer-rato, Lisa and Munson, Matthew and Dee, Stella and Gessner, Annette and Robertson, Bruce and Franzini, Greta and Stoyanova, Simona and Selle, Tabea and Springmann, Uwe and Clérice, Thibault. 2021a. *Corpus Scriptorum Ecclesiasticorum Latinorum digitized by Open Greek And Latin*.
- Crane, Gregory Ralph and Mylonas, Elli and Smith, Neel and Almas, Bridget and Babeu, Alison and Berra, Aurélien and Bonhomme, Marie-Laurence and Brooks, David and Cerrato, Lisa and Clérice, Thibault and Dee, Stella and Ehwald, R. and Frering, Léa and Gessner, An-nette and Harrington, J. Matthew and Himes, Zach and Jovanovic, Neven and Konieczny, Michael and Mahoney, Anne and Merrill, William and Mimno, David and Mueller, Lucian and Mun-son, Matthew and Singhal, Rashmi and Smith, David and Taounza-Jeminet, Alyx and Weaver, Gabriel A. and Wulfman, C. 2021b. *Perseus Canonical Latin Literature*.
- Anthonij Dees and Achim Stein and Pierre Kun-stmann and Martin Gleßgen and Giles Souvay. 2006. *Nouveau Corpus d'Amsterdam*.
- Draelants, Isabelle and Kuhry, Emmanuelle. 2020. *Sources des Encyclopédies Médiévale*.
- Glessgen, Martin-Dietrich and Carles, Hélène and Duval, Frédéric and Videsott, Paul. 2016. *Docu-ments linguistiques galloromans*.
- Goetz, Georg and Gitner, Adam, et al. *CGLO*. 23-12.
- Guillot, Céline and Heiden, Serge and Lavrentiev, Alexei. 2022. *Base de français médiéval*.
- Hasse, Dag Nikolaus. 2023. *Ptolemaeus Arabus et Latinus*.
- Neven Jovanović. 2024. *Croatiae Auctores Latini Textus*. Zenodo.
- Hannah Morcos and Simon Gaunt and Maria Teresa Rchetta and Henry Ravenhall and Si-mone Ventura and Luca Barbieri and Paul Caton and Geoffroy Noël and Ginestra Ferraro and Natasha Romanova. 2020. *The Histoire ancienne jusqu'à César: A Digital Edition (The Values of French)*.
- Perreaux, Nicolas. 2021. *Cartae Europae Medii Aevi (CEMA)*.
- Corinne Pierreville and Jules Nuguet and Ariane Pinche and Alexei Lavrentiev. 2025. *Projet Fa-bliaux*.
- Pinelli, Luca and Ctibor, Jan and Roelli, Philipp and Degl'Innocenti, Antonella and Gamberini, Roberto and Carlamaria Crespi, Serena and Fizzarotti, Luisa and Montepaone, Olivia and Vangone, Laura and Santarelli, Riccardo. 2023. *Mirabile Digital Archive*.
- Roelli, Philipp and Ctibor, Jan. 2024. *Corpus Cor-porum*.
- Russo, Emilio. 2023. *Biblioteca Italiana*.
- Schibel, Wolfgang and Kredel, Heinz and Niehl, Rüdiger. 1999–2008. *CAMENA: Latin Texts of Early Modern Europe*. DFG-funded project; on-line access.
- Clive R. Sneddon. 1982. *Old French corpus*. Liter-ary and Linguistic Data Service.
- Francesco Stella and Edoardo D'Angelo and Gior-gio Di Maria and Marina Buzzoni and Fulvio Delle Donne and Roberto Rosselli Del Turco. 2025. *ALIM: Archivio della Latinità Italiana del Medioevo*. Accessed: 2025-09-27.
- Stutzmann, Dominique and Torres Aguilar, Ser-gio and Chaffenet, Paul. 2021. *HOME-Alcar: Aligned and Annotated Cartularies*. Zenodo.
- Tabacco, Raffaella and Lana, Maurizio and Balossino, Michele and Bessi, Giancarlo and Bognini, Filippo and Buffa, Martina and Caso, Daniela and Cattaneo, Gianmario and Ciotti, Fabio and Ciusani, Mauro and Colombo, Piero and Cuzzotti, Claudia and del Core, Vincenzo and Della Calce, Elisa and Denicola, Luciano and Digirolamo, Letizia and Ferrandi, Etienne and Ferroni, Manuela and Fontana, Davide and Guiglia, Valeria and Loberti, Martina and Lucciano, Melanie and Maconi, Ludovica and Malaspina, Ermanno and Manuela, Ferroni and Marini, Alessia and Maronet, Léa and Massano, Federico and Mazzucco, Clementina and Mel-lano, Anastasia and Miglietta, Chiara and Mol-lea, Simone and Mosca, Laura and Musso, Si-mona and Naso, Manuela and Paniagua, David and Poncina, Fabio and Ramires, Giuseppe and Rinaldi, Valentina and Rosso, Nadia and Rota, Simona and Rozzi, Stefano and Rugnone, Elisa and Senore, Corinna and Stok, Fabio and Strona, Beatrice and Verny, Romain and Vitto-ria Martino, Maria. 2021. *Digital Library of Late-Antique Latin Texts*.
- Théry, Julien. 2022. *APOSCRIPTA database. Uni-fied Corpus of Papal Letters*. Zenodo.

A. Dataset and models

The ByT5 model can be found on HuggingFace at <https://huggingface.co/comma-project/normalization-byt5-small>.

The dataset is available on HuggingFace at <https://huggingface.co/datasets/comma-project/pen-alignment-pairs>.

B. Corpus Used for the Alignment

Corpus (Latin then French)	#kTokens
Corpus Corporum (div.) (Roelli and Ctibor, 2024)	96,995
CEMA (Perreux, 2021)	57,601
CAMENA (Schibel et al., 1999–2008)	7,916
(Assmann and Sahle, 1826–)	6,930
DigilibLT (Tabacco et al., 2021)	6,674
Perseus (Crane et al., 2021b)	6,148
CSEL (Crane et al., 2021a)	6,046
CroaLa (Jovanović, 2024)	5,233
Mirabile (Pinelli et al., 2023)	4,655
Bib. Italiana (Russo, 2023)	4,520
ALIM (Stella et al., 2025)	4,284
Add. Texts (Clérice, 2021)	3,844
Scripta (Bauduin et al., 2023)	2,226
Aposcripta (Théry, 2022)	1,736
ALCAR (Stutzmann et al., 2021)	973
Dig. Ptolemaeus (Hasse, 2023)	880
CGLO (Goetz and Gitner)	855
SourceEncyMe (Draelants and Kuhry, 2020)	777
NCA Dees et al. (2006)	3,185
BFM (Guillot et al., 2022)	6,293
Fabliaux (Pierreville et al., 2025)	120
Geste (Camps et al., 2019)	324
DocLing (Glessgen et al., 2016)	1,278
TVOF (Morcos et al., 2020)	361
OFC (Sneddon, 1982)	14
Total	229,870

Table B.5: Normalized corpus used by CoMMA for the ModernBert model.

C. Ranking of Corpora according to the Number of Alignments

Corpus	#alignments
Patrologia Latina	2,838,542
Lasciva Roma	413,580
CSEL	377,601
PerseusDL	306,297
DigilibLT	179,899
Auctores Scientiarum Varii	110,535
CEMA-latin	78,570

Table C.6: The most frequent aligned corpora in the dataset.

D. Llama3 8B Prompt

system: You are a specialist in Medieval Paleography and HTR correction.

Your goal is to transform "raw" HTR transcriptions into normalized Medieval Old French or Latin.

LINGUISTIC & PALAEOGRAPHIC RULES:

- HTR ERROR CORRECTION: Fix character confusions typical of HTR, such as 'm' being misread for 'in' (e.g., 'mexper' -> 'inexper') or 'f' for 's'.
- U/V & I/J NORMALIZATION: Distinguish between vocalic and consonantal uses. Use 'v' and 'j' for consonants and 'u' and 'i' for vowels.
- EXPANSION: Fully expand all scribal abbreviations and suspensions into their full character forms.
- GRAPHICAL VARIATION: Fully respect graphical variation, such as *filosofia* should be kept instead of normalizing to *philosophia*.
- APOSTROPHE: Add apostrophe where necessary (e.g., 'dun' -> 'd'un').
- SPACES: Remove extra spaces (e.g., 'coeur' -> 'coeur') and add necessary ones (e.g., 'letriste' -> 'le triste').

ABBREVIATION KEY:

- 9 or 9 = com / con (e.g., 9me -> comme)
- ʷ = us (e.g., voʷ -> vous)
- q̇ = qui
- † = ol (e.g., m†t -> molt)
- ̃ (tilde/macron) = nasalization (n/m) or missing letters (e.g., q̃ -> que)
- vertical tilde for suspension (missing letters) or R based phonemes (e.g., -re/er/-ra/ar-)

Return ONLY JSON.

user: Normalize this HTR output: text

assistant: