

From OCR to Analysis: Tracking Correction Provenance in Digital Humanities Pipelines

Haoze Guo

University of Wisconsin - Madison
College of Engineering
hguo246@wisc.edu

Ziqi Wei

University of Wisconsin - Madison
College of Letters and Sciences
zwei232@wisc.edu

Abstract

Optical Character Recognition (OCR) is a critical but error-prone stage in digital humanities text pipelines. While OCR correction improves usability for downstream NLP tasks, common workflows often overwrite intermediate decisions, obscuring how textual transformations affect scholarly interpretation. We present a provenance-aware framework for OCR-corrected humanities corpora that records correction lineage as base-anchored span edits, including edit type, correction source, confidence, and review status. Using a pilot corpus of historical texts, we compare downstream named entity extraction across raw OCR, fully corrected text, and provenance-filtered corrections. Our results show that correction pathways can substantially alter extracted entities and document-level interpretations, while provenance signals help identify unstable outputs and prioritize human review. We argue that provenance should be treated as a first-class analytical layer in NLP for digital humanities, supporting reproducibility, source criticism, and uncertainty-aware interpretation.

1 Introduction

Optical Character Recognition (OCR) is an important part of the digital humanities (DH) text analysis workflow which allows researchers to build searchable and analyzable corpora from scanned material (Smith, 2007; Nguyen et al., 2021). OCR is the first step in constructing resource-wide indices and retrieval systems for archival and newspaper collections and provides input for corpus-scale analysis processing pipelines (Ehrmann et al., 2020a; Düring et al., 2021, 2024). Unfortunately, OCR output produced from historical sources is often noisy because many historical materials are degraded, have non-standard characters, have unusual layouts, and have typographically varied characters (Nguyen et al., 2021; Neudecker et al., 2021). Consequently, many researchers use one or more

methods to clean up their OCR output prior to performing natural language processing (NLP) such as using normalizations through rules, post-correction through neural networks, and/or manually editing (Evershed and Fitch, 2014; Nguyen et al., 2020; Lyu et al., 2021).

These correction steps improve readability and frequently positively impact downstream NLP performance. However, they can also make it difficult to understand how the original text may have evolved throughout the process and what changes are still uncertain. Corrections can erase analytical history from a text by replacing the original OCR output without any indication of the original content, the nature of the change(s), or which change(s) are uncertain to the analyst. This case represents a methodological challenge for Digital Humanities (DH) because future interpretations of data (e.g., named entities, dates, topics, and trends) will be based on textual transformations that cannot be seen by the analyst (Strange et al., 2014; van Strien et al., 2020). This issue is particularly important for historical named entity recognition (NER) because noisy OCR and historical variation are known to have impacted the quality, reliability, and the ability to evaluate the data extracted from text (Hamdi et al., 2019; Ehrmann et al., 2024; Boroş et al., 2020).

We suggest that OCR correction be not just one way; but should be modeled as a repeated sequence of traceable editorial decisions; the provenance of which are also able to be examined with downstream NLP outputs. In the DH context this framing is especially important since core DH values (Ockeloen et al., 2013; Romein et al., 2024) include source criticism, uncertainty and interpretive transparency. Previous DH infrastructure work (Ockeloen et al., 2013; Guo and Wei, 2026a) has already emphasised the need for multi-level provenance as a significant requirement of scholarly workflows; therefore, broader provenance standards provide a

formal representation of such traces (Moreau et al., 2013; Lebo et al., 2013; Moreau and Groth, 2013). This context can also be seen by looking at broader audit-focused research that discusses the need for preserving transformation traces for longitudinal analysis and assessment (Guo and Wei, 2026c; Guo, 2026; Guo and Wei, 2026b), even across application domains.

In this paper, we introduce a provenance-aware representation for OCR-corrected humanities corpora and demonstrate its utility in a pilot study. We make three contributions:

- We introduce a span-level provenance schema for OCR correction that records edit lineage, correction source, confidence, and review status.
- We present a pilot empirical comparison of downstream named entity recognition (NER) across raw OCR, fully corrected text, and provenance-filtered corrected text.
- We provide a DH-oriented error analysis lens showing how provenance signals can identify unstable outputs and prioritize human review.

2 Related Work

Historical OCR pipelines face challenges beyond standard contemporary OCR settings. Noise can arise from page damage, bleed-through, historical fonts, line segmentation errors, footnotes, marginalia, and spelling conventions that differ from modern forms (Nguyen et al., 2021; Neudecker et al., 2021). As a result, OCR errors are not uniformly distributed across documents or regions of a page, and “correction” may involve both error repair and editorial normalization (Evershed and Fitch, 2014; Nguyen et al., 2021).

For DH scholars, this distinction matters. A correction that restores a malformed token may improve fidelity to the source, while a normalization step may improve computational consistency but alter historically meaningful variation (Strange et al., 2014). In both cases, downstream NLP outputs can shift: entities may appear or disappear, names may merge or split, and frequencies may change (van Strien et al., 2020; Hamdi et al., 2019; Ehrmann et al., 2024). If the intermediate correction decisions are not preserved, it becomes difficult to audit why a downstream result changed.

This issue is especially salient for named entity analysis on historical texts. Shared tasks and survey work (e.g., HIPE) document the difficulty

of robust NER and linking in multilingual historical materials, including challenges from noisy OCR, diachronic variation, and annotation heterogeneity (Ehrmann et al., 2020b, 2022, 2024). Parallel work on historical newspaper infrastructures such as *impresso* has also emphasized the need to connect computational processing with scholarly interpretability and interface-level transparency (Ehrmann et al., 2020a; Düring et al., 2021, 2024).

Our motivation is therefore not only to improve OCR correction quality, but to preserve and expose the lineage of correction decisions so that computational findings remain inspectable and reproducible. Provenance research offers a strong conceptual foundation for this goal, including explicit modeling of entities, activities, and agents (Moreau et al., 2013; Lebo et al., 2013), but DH OCR pipelines need a lightweight, text-first representation that can operate at the span level and integrate directly with downstream NLP outputs.

We summarize how common provenance/annotation and OCR-output practices relate to the requirements of replayable span edits and policy-driven variant construction in Table 1. This comparison motivates our focus on base-anchored span-edit events coupled with explicit trust policies for reconstruction and audit.

3 Provenance-Aware Correction Schema

3.1 Design Principles

We design the schema around three principles: (1) **traceability**, so each correction can be linked to a document location and prior text span; (2) **tool-agnostic interoperability**, so the representation can capture rule-based, model-based, and human edits; and (3) **analytical usefulness**, so metadata can support downstream filtering and audit. Our design is informed by both DH provenance modeling practices and the broader PROV family of standards (Ockeloen et al., 2013; Moreau et al., 2013; Lebo et al., 2013).

We represent corrections at the **span level**, which provides more flexibility than token-only records while remaining simpler than character-level edit traces. A span may correspond to a token, multi-token phrase, or split/merge operation. This design is compatible with OCR post-correction workflows ranging from noisy-channel approaches to neural sequence models (Evershed and Fitch, 2014; Nguyen et al., 2020; Lyu et al., 2021; Nguyen et al.,

Approach	Span-level edits	Base-anchored offsets	Deterministic replay	Trust-policy filtering	Downstream trace links	Pipeline-ready artifact
PROV-DM / PROV-O (Moreau et al., 2013; Lebo et al., 2013)	✗	✗	✗	✗	✓	✗
TEI stand-off (<standoff>) (Text Encoding Initiative, 2026)	✓	✓	✗	✗	✗	✓
ALTO XML (Library of Congress, 2022)	✗	✓	✗	✗	✗	✓
hOCR (Bohnet et al., 2024)	✗	✓	✗	✗	✗	✓
PAGE-XML (Pletschacher and Antonacopoulos, 2010)	✗	✓	✗	✗	✗	✓
OCR-D (Neudecker et al., 2019; OCR-D Consortium, 2026b,a)	✗	✓	✗	✗	✗	✓
Span-edit events + provenance-aware filtering (Ours)	✓	✓	✓	✓	✓	✓

Table 1: Comparison of our span-edit provenance + policy-driven reconstruction with common DH/OCR provenance and annotation practices.

2021).

3.2 Schema

We represent OCR correction provenance at the span level using a compact record schema that links each edit to a document location and revision step. Each record stores: document/page identifiers, span offsets, original and corrected text, edit type (e.g., substitution, split, merge), correction source (rule-based, model-assisted, or human), optional confidence, optional review status, and optional layout-zone metadata (e.g., body, header, footnote). This representation is sufficient to reconstruct correction lineage and to trace downstream NLP outputs (e.g., extracted entities) back to the edits they depend on (Ciccarese et al., 2013).

3.3 Illustrative Record

A single provenance record links an OCR span to its corrected form together with the metadata needed for audit (e.g., edit type, correction source, confidence, and review status). For example, if the OCR span *Madifon* is corrected to *Madison* by a model-assisted step with confidence 0.74 and no human approval, any downstream entity extracted from *Madison* can be traced back to that specific correction event. This span-level linkage is central to our analysis of entity volatility and uncertainty-aware filtering in the pilot study.

3.4 Serialization and Interoperability

Our span-level provenance schema is designed to be tool-agnostic and easy to integrate into existing NLP pipelines. The correction records can be serialized as **JSONL**, as a **tabular file** (CSV/Parquet), or as **stand-off annotations** that reference the underlying text by character offsets. Stand-off serialization is particularly useful for DH settings because it preserves the original OCR text while enabling multiple correction layers to be applied or compared without overwriting intermediate states.

In JSON-based NLP pipelines, provenance records can be embedded as an auxiliary field

alongside the raw and corrected text variants, enabling downstream tasks to (i) reconstruct a chosen variant under a specified trust policy, and (ii) trace extracted entities back to the correction events that influenced them. For interoperability across tools, we recommend storing: stable document/page identifiers, explicit span offsets, explicit revision steps, and correction-source metadata, so that records remain portable even when tokenization or sentence segmentation changes. This makes the schema compatible with common DH workflows and with standard NLP processing stacks that expect JSON-serializable artifacts.

3.5 Offset Semantics and Edit Application

All correction events are anchored to a **base revision** (typically raw OCR) to avoid cascading offset drift and to preserve stable linkage from downstream outputs back to scan-derived text. Offsets (span_start, span_end) are Unicode-codepoint indices over the base text using half-open intervals [start,end); orig_text must exactly match the base substring at that interval as an integrity check. To construct a variant, we select events under a trust policy (e.g., confidence threshold or review_status=approved), sort by span_start (tie-break by event_id), and apply them as base-anchored replacements, explicitly detecting overlaps for resolution (e.g., prefer human over model or defer to adjudication). Boundary-sensitive edits (split/merge) and OCR normalizations (line-break/whitespace reflow, hyphenation repair such as "inter-national"→"international", paragraph merges) are represented as the same span-replacement events with explicit boundaries; large reflow regions are recorded as a single replacement to avoid brittle micro-edits.

4 Pilot Study Design

4.1 Corpus and Unit of Analysis

We conduct a pilot study on a small corpus of historical texts drawn from a digitized humanities collec-

tion. The corpus includes scanned materials with OCR output and a subset of passages that undergo correction. We treat this study as a methodological pilot rather than a benchmark-scale evaluation, with the goal of testing whether provenance signals provide useful analytical leverage in downstream NLP (van Strien et al., 2020).

Our primary unit of analysis is the **document-level NER output** for each text variant, with additional linkage to **span-level correction records**. For each document, we preserve the raw OCR text and record every applied correction as a provenance event, including span offsets, edit type, correction source, and optional confidence/review status. This design supports both aggregate comparison, how the extracted entity inventory changes across variants, and auditability which correction events are associated with unstable entities.

4.2 Text Variants and Provenance Filtering Policy

For each document, we construct three text variants from the same OCR source: **Raw OCR** (no post-correction), **Fully corrected** (all available corrections applied), and **Provenance-filtered** (only corrections meeting a provenance criterion).

The provenance-filtered variant operationalizes a conservative “trust policy” over corrections. In the main comparison, this policy is confidence-based (e.g., confidence ≥ 0.70) and may optionally require human approval when such metadata is available. We additionally report a threshold sensitivity analysis (Table 3) that varies the filter strictness, including higher confidence thresholds and a human-approved-only condition. This provides a transparent characterization of the coverage–stability tradeoff induced by provenance-aware filtering, rather than presenting a single corrected corpus as an implicit ground truth.

This design separates two pipeline questions that are often conflated in DH OCR workflows: (i) whether correction improves downstream extraction at all, and (ii) which corrections should be treated as sufficiently reliable for interpretive analysis.

4.2.1 Where Confidence Comes From

In our pilot, confidence is an optional provenance attribute attached to each correction event and emitted by the correction pathway that proposed the edit. For **model-assisted** corrections, confidence corresponds to the model’s own edit score, used to rank

suggested substitutions and split/merge operations. For **rule-based** corrections, confidence is derived from rule certainty (e.g., high-confidence deterministic patterns versus lower-confidence heuristic matches). For **human** corrections, we treat `review_status=approved` as the primary reliability signal; when an explicit reviewer score is available, we store it as confidence.

These scores are not assumed to be calibrated probabilities, and they are not assumed comparable across different tools or correction sources. Accordingly, we use confidence only as a *within-pipeline ranking signal* for sensitivity analysis. Thresholds (e.g., confidence ≥ 0.70 in the main provenance-filtered condition) are applied under a fixed correction stack and interpreted operationally rather than as a universal notion of correctness. The threshold sweep (Table 3) is therefore intended to make the coverage–stability tradeoff explicit under a fixed scoring regime, not to claim that a given numeric threshold generalizes across tool stacks.

4.3 Downstream Task and Quantitative Metrics

We use named entity recognition (NER) as the downstream task because entity extraction is common in DH workflows and is highly sensitive to OCR noise, normalization, and segmentation decisions (Hamdi et al., 2019; Boros et al., 2020; Ehrmann et al., 2020b, 2022, 2024; Clausner et al., 2011). We run the same NER pipeline over all text variants and compare the resulting entity outputs.

4.3.1 NER Pipeline Configuration

We use a single, fixed NER pipeline across all text variants to isolate the effect of OCR correction and provenance filtering. Concretely, we sentence-segment the text, tokenize using the model’s native subword tokenizer, and apply a transformer-based NER model fine-tuned on CoNLL-2003 (implemented via the HuggingFace Transformers library) (Devlin et al., 2019; Liu et al., 2019; Wolf et al., 2020). We keep all inference settings fixed across variants.

Because historical corpora exhibit diachronic spelling and orthographic variation, we avoid additional spelling modernization or normalization inside the NER pipeline beyond what is already implied by the OCR correction variants under study. This design prevents conflating “downstream robustness tricks” with the provenance-aware correction analysis. We acknowledge that domain

adaptation can substantially affect absolute NER accuracy in historical settings; our goal here is to characterize how correction pathways change extracted entities under a consistent downstream model (Ehrmann et al., 2024; Hamdi et al., 2019; Boros et al., 2020).

To quantify differences across variants, we report:

- **Entity mentions:** total number of extracted entity mentions per variant.
- **Unique entities:** number of distinct extracted entity surface strings per variant.
- **Entity overlap:** Jaccard similarity between unique-entity sets across variants.
- **Entity volatility:** entities that (i) appear in one variant but not another, or (ii) appear in both but with changed surface form or span boundary.

In addition, we compute the share of volatile entities linked to **low-confidence or unreviewed correction spans**. This measure provides an audit-oriented view of instability: instead of treating volatility as an opaque downstream artifact, it tests whether provenance fields can identify subsets of edits that disproportionately contribute to unstable outputs.

4.4 Entity Linking as a Secondary Downstream Task

To test whether provenance-aware correction affects not only entity extraction but also downstream *resolution* of entities to knowledge bases, we include a lightweight entity linking (EL) analysis on a pilot subset. For each text variant, we take extracted entity mentions and apply an off-the-shelf EL system to link mentions to a reference knowledge base (e.g., Wikipedia/Wikidata identifiers). We focus on mentions that are volatile across variants, since these are the cases where small surface-form or boundary changes are most likely to alter linking decisions.

We report (i) **linking coverage** (fraction of mentions assigned a KB identifier), (ii) **link stability** across variants, fraction of mentions that resolve to the same KB identifier, and (iii) a small **manual audit of disagreements** to characterize whether changes reflect improved disambiguation, harmful normalization, or overconfident linking on noisy strings. This EL check is designed as a methodological stress test rather than a benchmark-scale evaluation.

4.5 Attribution of Entity Changes to Correction Events

To associate downstream entity differences with provenance records, we compute an **association** (not causal proof) between each entity mention and nearby correction events. We first attempt **span overlap**: an entity mention is associated with a correction record if the mention character offsets overlap the correction span offsets in that same variant. For edits that may shift offsets (notably split/merge operations), we use a bounded local fallback: we search for correction events within a window of $\pm W$ characters around the entity mention (we set $W = 50$ in the pilot) and choose the closest event by absolute offset distance. Ties are broken by preferring (i) overlap over proximity, (ii) edit types that can directly affect boundaries (split/merge) over substitutions, and (iii) higher-confidence events when available.

We use this association to support audit and summarization (e.g., “how many volatile entities are linked to low-confidence edits”), and we interpret results as identifying *likely contributing edits* rather than definitive causes.

4.6 Validation of the Attribution Heuristic

Because proximity-based association can produce spurious links (nearby edits may be unrelated), we validate the attribution heuristic on a labeled sample. We randomly sample linked (entity, correction-event) pairs and ask annotators to judge whether the correction plausibly contributed to the entity change (yes/no). We report the precision of the heuristic and use this validation to qualify interpretation of the signal-utility analysis.

4.7 Qualitative Coding Protocol

To complement the quantitative comparison, we manually inspect a pilot sample of volatile entities and assign each case to an error category (e.g., OCR noise, normalization shift, split/merge boundary issue, layout artifact). The goal of this coding step is not to produce an exhaustive taxonomy, but to evaluate whether provenance fields (confidence, edit type, layout zone, review status) provide actionable diagnostic value. We use the qualitative categories to interpret the mechanisms behind volatility and to assess which provenance signals are most useful for DH-oriented review and source criticism.

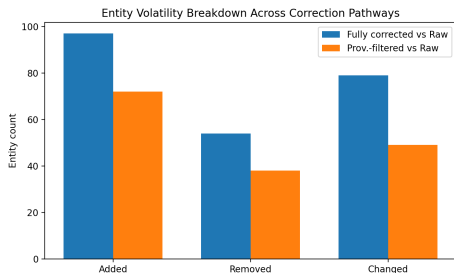


Figure 1: Breakdown of entity volatility across correction pathways.

5 Results

5.1 Quantitative Differences Across Correction Pathways

Across the pilot corpus, correction activity is unevenly distributed across documents and page regions, and many edits occur in short spans that affect candidate entity strings. This matters because a relatively small number of span edits can produce disproportionately large downstream changes in NER outputs, particularly when edits touch capitalization, rare proper nouns, or token boundaries.

Figure 1 provides a compact view of entity volatility relative to raw OCR. The fully corrected condition produces more *added* and *changed* entities than the provenance-filtered condition, indicating stronger transformation of the extracted entity inventory. The provenance-filtered condition still preserves a substantial portion of correction gains, but reduces high-risk changes. In practice, this suggests that provenance filtering behaves less like “undoing correction” and more like selecting a correction pathway with a different analytical risk profile.

Table 2 quantifies these differences. Relative to raw OCR, the fully corrected variant increases both entity mentions (1184 \rightarrow 1342) and unique entities (512 \rightarrow 566), consistent with correction improving recognizability of names and places. However, these gains coincide with the highest volatility (176 volatile entities), meaning that correction alters not only the volume of extracted entities but also the stability of the extracted inventory. The provenance-filtered variant retains most of the coverage gains (1287 mentions; 548 unique entities) while reducing volatility (121), suggesting that provenance-aware filtering can improve analytical stability without reverting to raw OCR.

A key observation is that volatility is not uniformly distributed across edits: a substantial frac-

Variant	Ment.	Uniq.	Jac. vs Raw	Vol.	% Unrev.
Raw OCR	1184	512	1.00	–	–
Fully corrected	1342	566	0.69	176	68%
Prov.-filtered	1287	548	0.76	121	76%

Table 2: NER output differences across OCR correction pathways. “Vol.” counts entities that appear/disappear or change surface form/span boundary across variants; “% Unrev.” is the share of volatile entities linked to unreviewed edits.

tion of volatile entities in the corrected conditions is linked to low-confidence or unreviewed correction events (Table 2). This does not imply that all low-confidence edits are incorrect; rather, it indicates that these edits are most likely to produce downstream differences that analysts may want to inspect, qualify, or report explicitly.

5.2 Threshold Sensitivity of Provenance Filtering

Table 3 shows how downstream outputs change as the provenance filter becomes stricter. As the confidence threshold increases (or when only human-approved edits are used), entity volatility decreases monotonically, but entity coverage also declines. This tradeoff is expected and analytically useful: different DH workflows may prefer different operating points depending on whether the priority is recall-oriented exploration or conservative interpretation.

Figure 2 visualizes the same sweep as a coverage–stability curve. Moving from “All corrections” to stricter thresholds shifts the operating point toward lower instability, but also lower coverage. In this pilot, a mid-range threshold (e.g., confidence ≥ 0.70) yields a balanced regime: many of the coverage improvements from correction are retained, while a substantial portion of volatility is reduced. We emphasize that thresholds are not universal: confidence values are tool- and workflow-dependent. The contribution here is that provenance makes the operating point explicit and reportable, rather than implicit in a single overwritten corrected text.

5.3 Provenance Signals Predicting Instability

Beyond overall thresholds, provenance fields help identify which correction operations are most likely to affect downstream entity extraction. Table 4 summarizes the utility of several provenance signals for predicting volatility. Two patterns are par-

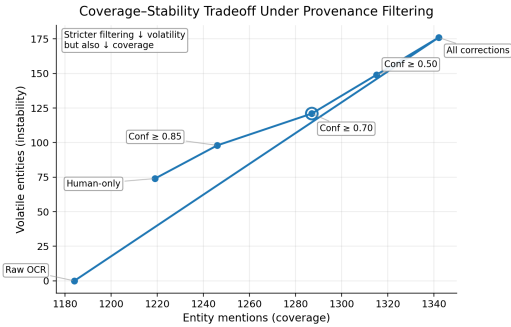


Figure 2: Coverage–stability tradeoff under provenance filtering.

Filter	Mentions	Unique	Jaccard vs Raw	Volatile
Raw OCR	1184	512	1.00	–
All corrections	1342	566	0.69	176
Conf \geq 0.50	1315	559	0.73	149
Conf \geq 0.70	1287	548	0.76	121
Conf \geq 0.85	1246	531	0.79	98
Human-approved only	1219	523	0.82	74

Table 3: Threshold sensitivity for provenance-filtered correction.

ticularly salient. First, boundary-affecting edits (split/merge) exhibit the highest volatility lift despite being relatively infrequent, consistent with entity extraction being sensitive to segmentation decisions. Second, non-body layout zones (e.g., headers and footnotes) are instability hotspots, reflecting known OCR challenges in regions where typography and layout differ from running prose.

Low confidence and unreviewed status provide coarse but useful uncertainty flags. These fields are less diagnostic than edit type or layout zone, but they are broadly applicable across correction pathways and can support conservative analysis regimes (e.g., confidence-based filtering) and targeted inspection when human review capacity is limited.

These associations are produced by a deterministic overlap-plus-window heuristic and should be interpreted as likely contributing edits rather than definitive causal attributions; we therefore validate heuristic precision on a labeled sample and report it alongside the signal-utility analysis.

5.4 Qualitative Error Categories and Diagnostic Use

Table 5 summarizes qualitative categories observed among volatile entities. The most common category is false entities caused by OCR noise, fol-

Signal	%	Flag vol.	Lift
Low conf. ($<$ 0.70)	31%	7.8%	2.7 \times
Unreviewed	90%	5.9%	1.7 \times
Split/merge	9%	12.5%	3.3 \times
Non-body zone	14%	10.1%	2.6 \times

Table 4: Utility of provenance signals for identifying instability. “Flag vol.” is the volatility rate among entities associated with the flagged condition; “Lift” compares flagged vs unflagged.

lowed by normalization shifts and boundary errors (split/merge cases). These patterns align with known OCR and historical NER challenges (van Strien et al., 2020; Ehrmann et al., 2024), but the key result in our setting is that provenance metadata helps localize and explain failures rather than only reporting aggregate quality.

Different provenance fields are useful for different diagnostic purposes. Confidence and review status help identify risky corrections for triage; edit type helps distinguish normalization shifts from segmentation effects; and layout-zone metadata supports page-region-based auditing of headers, footnotes, and other non-body text. Taken together, provenance enables analysts to answer not only “what changed” (entity volatility), but also “which editorial operations caused the change” and “which changes are plausible candidates for targeted review.”

5.5 Entity Linking Sensitivity to Correction Pathways

Entity linking is particularly sensitive to surface form and boundary decisions: small spelling changes, whitespace insertion, or normalization can alter candidate generation and disambiguation. In our pilot EL analysis, fully corrected text increases linking coverage relative to raw OCR by producing more linkable surface strings, but it also introduces additional link instability for a subset of volatile mentions. Provenance-filtered text typically preserves much of the coverage gain while reducing the rate of link changes tied to low-confidence or unreviewed corrections.

Qualitatively, we observe three recurring EL disagreement patterns: (i) **repair-driven improvements** where correcting OCR noise enables linking to the intended entity, (ii) **normalization-driven shifts** where historical spellings or abbreviations are modernized and shift candidate selection, and (iii) **boundary-driven failures** where split/merge

Pattern	%	Diagnostic provenance cue(s)
OCR-noise false entity	29%	Low confidence; model-assisted; often disappears under filtering
Normalization shift	21%	edit_type=substitution; unreviewed; surface-form changes
Boundary (split/merge)	18%	edit_type=split/merge; span boundary drift
Layout artifact	17%	layout_zone=header/footer; non-body text
Ambiguous correction	15%	Medium confidence; unreviewed; competing entity strings

Table 5: Qualitative categories among volatile entities.

edits change mention spans and trigger different candidate sets. These results reinforce the core claim that correction pathways are analytically consequential: downstream interpretations can shift not only in *which* entities are extracted, but also in *which real-world entities* those mentions are resolved to.

6 Discussion

The pilot study shows that OCR correction pathways can produce meaningfully different analytical corpora from the same scanned source. Raw OCR, fully corrected text, and provenance-filtered text do not merely differ in readability; they produce different entity inventories, different levels of volatility, and different evidentiary risks for downstream interpretation. This supports our central claim that OCR correction should be treated as an inspectable analytical process rather than a hidden preprocessing step.

Provenance makes these differences explicit. By preserving edit source, confidence, review status, edit type, and layout-zone information, analysts can treat correction as a set of documented editorial choices rather than as an invisible transformation. The threshold sweep and coverage–stability curve show how different trust policies expose different operating points: a recall-oriented analysis may prefer broader correction coverage, while a conservative interpretive analysis may prefer provenance-filtered text with lower volatility.

The qualitative categories highlight a recurring DH concern: some corrections restore fidelity to the source (repair), while others normalize historically meaningful variation (normalization). Provenance makes these transformations visible and therefore contestable. This is important for interpretive work where spelling variation, abbreviations, and typography may carry historical meaning.

Provenance also supports locating where instability originates. Our signal-utility analysis suggests that boundary edits and non-body layout zones are particularly consequential. This enables a more DH-aligned auditing practice in which analysts can attribute changes in extracted entities to specific editorial operations, rather than treating entity differences as opaque model behavior.

6.1 Limitations and Future Work

This study has several limitations. First, it is a methodological pilot rather than a benchmark-scale evaluation, and the corpus is not intended to represent the full diversity of DH materials. Future work should test the schema across multiple DH domains, including historical newspapers, literary prose, poetry, correspondence, and archival administrative records, because each genre introduces different OCR and correction failure modes. Second, we use NER as the primary downstream probe; additional tasks such as topic modeling, text classification, relation extraction, and entity linking at larger scale may expose different kinds of correction sensitivity. Additionally, although confidence scores and metadata can be utilized with some degree of familiarity to compare various corrections, they serve as a signal of provenance rather than as calibrated probabilities.

Future research into the analysis of correction pathways which account for provenance must be extended to cover multilingual collections and more complex historical page layouts. In these instances, greater differences in typography and segmentation between different collections will be present. A second area for further research is to expand the downstream tasks used for example with regard to entity linking where small changes in the surface (visibility) features result in significant impact on linking decisions. Lastly, the construction of provenance-aware pipelines creates an opportunity to establish standardized reporting criteria for DH-related NLP. Researchers will be able to use these documents as a means of reporting their own sensitivity to correction pathways while also documenting the provenance signals that directly contributed most to their instability.

7 Conclusion

We presented a provenance-aware framework for OCR-corrected humanities corpora and a pilot study showing how correction pathways alter down-

stream named entity extraction. By preserving correction lineage at the span level, our approach helps make NLP outputs more auditable, reproducible, and interpretable in DH workflows. We argue that provenance should be treated as a first-class analytical layer in OCR-to-NLP pipelines, not merely as implementation metadata.

References

- Klaus Bohnet and 1 others. 2024. [hocr: The embedded ocr workflow and output format \(specification\)](#). Specification (GitHub repository).
- Emanuela Boroş, Ahmed Hamdi, Elvys Linhares Pontes, Luis Adrián Cabrera-Diego, Jose G. Moreno, Nicolas Sidere, and Antoine Doucet. 2020. [Alleviating digitization errors in named entity recognition for historical documents](#). In *Proceedings of the 24th Conference on Computational Natural Language Learning (CoNLL)*, pages 431–441.
- Paolo Ciccarese, Stian Soiland-Reyes, Khalid Belhajjame, Alasdair J. G. Gray, Carole Goble, and Tim Clark. 2013. Pav ontology: Provenance, authoring and versioning. *Journal of Biomedical Semantics*.
- Christian Clausner, Stefan Pletschacher, and Apostolos Antonacopoulos. 2011. [Scenario driven in-depth performance evaluation of document layout analysis methods](#). In *Proceedings of the International Conference on Document Analysis and Recognition (ICDAR)*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of NAACL-HLT*.
- Marten Düring, Estelle Bunout, and Daniele Guido. 2024. [Transparent generosity: Introducing the impresso interface for the exploration of semantically enriched historical newspapers](#). *Historical Methods: A Journal of Quantitative and Interdisciplinary History*, 57(1):20–40.
- Marten Düring, Roman Kalyakin, Estelle Bunout, and Daniele Guido. 2021. [Impresso inspect and compare: Visual comparison of semantically enriched historical newspaper articles](#). *Information*, 12(9):348.
- Maud Ehrmann, Ahmed Hamdi, Elvys Linhares Pontes, Matteo Romanello, and Antoine Doucet. 2024. [Named entity recognition and classification in historical documents: A survey](#). *ACM Computing Surveys*, 56(2):27:1–27:47.
- Maud Ehrmann, Matteo Romanello, Simon Clematide, Phillip Benjamin Ströbel, and Raphaël Barman. 2020a. [Language resources for historical newspapers: The impresso collection](#). In *Proceedings of the 12th Language Resources and Evaluation Conference (LREC 2020)*, pages 958–968.
- Maud Ehrmann, Matteo Romanello, Alex Flückiger, and Simon Clematide. 2020b. [Overview of clef hipe 2020: Named entity recognition and linking on historical newspapers](#). In *Experimental IR Meets Multilinguality, Multimodality, and Interaction (CLEF 2020)*, volume 12260 of *Lecture Notes in Computer Science*, pages 288–310.
- Maud Ehrmann, Matteo Romanello, Sven Najem-Meyer, Antoine Doucet, and Simon Clematide. 2022. [Extended overview of hipe-2022: Named entity recognition and linking in multilingual historical documents](#). In *Working Notes of CLEF 2022*, volume 3180 of *CEUR Workshop Proceedings*, pages 1038–1063.
- John Evershed and Kent Fitch. 2014. [Correcting noisy ocr: Context beats confusion](#). In *Proceedings of the First International Conference on Digital Access to Textual Cultural Heritage (DATECH 2014)*, pages 45–51.
- Haoze Guo. 2026. [Consentdiff at scale: Longitudinal audits of web privacy policy changes and ui frictions](#). In *Proceedings of the Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems, CHI EA '26*, New York, NY, USA. Association for Computing Machinery.
- Haoze Guo and Ziqi Wei. 2026a. [Behind the feed: A taxonomy of user-facing cues for algorithmic transparency in social media](#). *Preprint*, arXiv:2602.03121.
- Haoze Guo and Ziqi Wei. 2026b. [Hidden-in-plain-text: A benchmark for social-web indirect prompt injection in rag](#). In *Proceedings of the ACM Web Conference 2026, WWW '26*, page 8337–8340, New York, NY, USA. Association for Computing Machinery.
- Haoze Guo and Ziqi Wei. 2026c. [Temporal drift in privacy recall: Users misremember from verbatim loss to gist-based overexposure](#). In *Proceedings of the 2025 International Conference on Human-Engaged Computing, ICHEC '25*, New York, NY, USA. Association for Computing Machinery.
- Ahmed Hamdi, Axel Jean-Caurant, Nicolas Sidere, Mickaël Coustaty, and Antoine Doucet. 2019. [An analysis of the performance of named entity recognition over ocred documents](#). In *2019 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*, pages 333–334.
- Timothy Lebo, Satya Sahoo, Deborah McGuinness, and 1 others. 2013. [Prov-o: The prov ontology](#). W3C Recommendation.
- Library of Congress. 2022. [Alto: Technical metadata for layout and text objects](#). Standard (Library of Congress).
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). In *arXiv*.

- Lijun Lyu, Maria Koutraki, Martin Krickl, and Besnik Fetahu. 2021. [Neural ocr post-hoc correction of historical corpora](#). *Transactions of the Association for Computational Linguistics*, 9:479–493.
- Luc Moreau and Paul Groth. 2013. [Provenance: An Introduction to PROV](#), volume 3 of *Synthesis Lectures on the Semantic Web: Theory and Technology*. Morgan & Claypool.
- Luc Moreau, Paul Groth, and 1 others. 2013. [Prov-dm: The prov data model](#). W3C Recommendation.
- Clemens Neudecker, Konstantin Baierer, Don Borgman, Markus Federbusch, Matthias Boenig, Wolfgang Maier, Christian Reul, and Frank Puppe. 2019. [Ocr-d: An end-to-end open source ocr framework for historical printed documents](#). In *Proceedings of the 3rd International Conference on Digital Access to Textual Cultural Heritage (DATECH 2019)*.
- Clemens Neudecker, Konstantin Baierer, Mike Gerber, Christian Clausner, Apostolos Antonacopoulos, and Stefan Pletschacher. 2021. [A survey of ocr evaluation tools and metrics](#). In *Proceedings of the 6th International Workshop on Historical Document Imaging and Processing (HIP 2021)*, pages 13–18.
- Thi Tuyet Hai Nguyen, Adam Jatowt, Mickaël Coustaty, and Antoine Doucet. 2021. [Survey of post-ocr processing approaches](#). *ACM Computing Surveys*, 54(6):124:1–124:37.
- Thi-Tuyet-Hai Nguyen, Adam Jatowt, Nhu-Van Nguyen, Mickaël Coustaty, and Antoine Doucet. 2020. [Neural machine translation with bert for post-ocr error detection and correction](#). In *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries (JCDL 2020)*, pages 333–336.
- C. J. Ockelo, A. S. Fokkens-Zwirello, S. ter Braake, P. T. J. M. Vossen, V. de Boer, A. T. Schreiber, and S. Legene. 2013. [Biographynet: Managing provenance at multiple levels and from different perspectives](#). In *Proceedings of the 3rd International Workshop on Linked Science (LISC 2013)*.
- OCR-D Consortium. 2026a. [Ocr-d workflow guide](#). OCR-D Documentation.
- OCR-D Consortium. 2026b. [Requirements on handling mets/page](#). OCR-D Specification.
- Stefan Pletschacher and Apostolos Antonacopoulos. 2010. [The page \(page analysis and ground-truth elements\) format framework](#). In *Proceedings of the 20th International Conference on Pattern Recognition (ICPR)*.
- C. Annemieke Romein, Tobias Hodel, Femke Gordijn, Joris van Zundert, Alix Chagué, Milan van Lange, Helle Strandgaard Jensen, Andy Stauder, Jake Purcell, Melissa Terras, Pauline van den Heuvel, Carlijn Keijzer, Achim Rabus, Chantal Sitaram, Aakriti Bhatia, Katrien Depuydt, Mary Aderonke Afolabi, Anastasiia Anikina, Elisa Bastianello, and 1 others. 2024. [Exploring data provenance in handwritten text recognition infrastructure: Sharing and reusing ground truth data, referencing models, and acknowledging contributions](#). *Journal of Data Mining and Digital Humanities*.
- Ray Smith. 2007. [An overview of the tesseract ocr engine](#). In *Proceedings of the Ninth International Conference on Document Analysis and Recognition (ICDAR 2007)*, pages 629–633.
- Carolyn Strange, Daniel McNamara, Josh Wodak, and Ian Wood. 2014. [Mining for the meanings of a murder: The impact of ocr quality on the use of digitized historical newspapers](#). *Digital Humanities Quarterly*, 8(1).
- Text Encoding Initiative. 2026. [The standOff element \(tei p5 guidelines\)](#). TEI Guidelines reference.
- Daniel van Strien, Kaspar Beelen, Mariona Coll Ardanuy, Kasra Hosseini, Barbara McGillivray, and Giovanni Colavizza. 2020. [Assessing the impact of ocr quality on downstream nlp tasks](#). In *Proceedings of the 12th International Conference on Agents and Artificial Intelligence (ICAART 2020)*, pages 484–496.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clément Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of EMNLP: System Demonstrations*.

A Schema Specification and Serialization

A.1 Normative field specification

We represent OCR correction provenance as a sequence of **span-edit events** anchored to a **base revision** of the text (typically raw OCR). Each event specifies *where* an edit applies in the base text, *what* replacement occurs, and *metadata* that supports audit and policy-based reconstruction.

Required invariants. Each event **MUST** satisfy *base anchoring* (offsets refer to `base_revision`, with `base_revision=0` for raw OCR in this paper), *half-open spans* (`[span_start, span_end)` with `span_start < span_end`), *Unicode-codepoint offsets*, and an *integrity check* (`orig_text` matches the base substring at `[span_start, span_end)`).

Event fields. Table 6 defines the canonical fields and constraints used by the schema described in §3 and referenced throughout the pilot study.

Versioning. We include `schema_version` in each event. Minor versions may add optional fields without changing existing field meanings; major versions may change field meanings or constraints.

Field	Type	Req.	Meaning / constraints
schema_version	string	Y	Semantic version, e.g., "1.0.0".
event_id	string	Y	Unique identifier within a corpus (UUID recommended).
doc_id	string	Y	Stable document identifier.
page_id	string/int	Y	Page identifier (or region id if pageless).
base_revision	int	Y	Revision index of anchor text (0 = raw OCR).
span_start	int	Y	Inclusive offset in base revision (Unicode codepoints).
span_end	int	Y	Exclusive offset; $\text{span_start} < \text{span_end}$.
orig_text	string	Y	Exact base substring at $[\text{span_start}, \text{span_end})$.
new_text	string	Y	Replacement string; empty string indicates deletion.
edit_type	enum	Y	{substitute, insert, delete, split, merge, normalize}.
source	enum	Y	{rule, model, human}.
confidence	float	O	In $[0,1]$; tool-specific ranking score (not assumed calibrated).
review_status	enum	O	{unreviewed, approved, rejected}.
reviewer_id	string	O	Pseudonymous reviewer id, if applicable.
layout_zone	enum/string	O	e.g., body, header, footnote, caption.
note	string	O	Free-form rationale/tag (e.g., "hyphenation repair").

Table 6: Canonical span-edit provenance event schema used in this paper.

All experiments in the current draft assume a fixed schema.

A.2 Canonical JSONL and OCR operation examples

We serialize events as JSONL (one event per line). This supports streaming pipelines and common NLP artifact formats while keeping events queryable and replayable.

Canonical example. The main text uses an illustrative correction *Madifon*→*Madison* (model-assisted; confidence 0.74; no human approval). In our schema this is represented as:

```
{ "schema_version": "1.0.0", "event_id": "c2f1...",
  "doc_id": "doc_017", "page_id": 3, "base_revision": 0,
  "span_start": 1284, "span_end": 1291,
  "orig_text": "Madifon", "new_text": "Madison",
  "edit_type": "substitute", "source": "model",
  "confidence": 0.74, "review_status": "unreviewed",
  "layout_zone": "body" }
```

OCR-specific operations as span edits. We encode OCR cleanup as span replacements: *hyphenation repair* ("international"→"international") via normalize or substitute spanning the hyphen and break; *line-break normalization* (newline→space) via normalize; *whitespace reflow* (collapse repeated whitespace / remove intra-word spaces) via normalize; *split/merge boundaries* (e.g., "NewYork"→"New York") via split/merge; and *paragraph merges/reflow* as a single larger-span replacement to avoid brittle micro-edits while preserving auditability.

B Provenance-Aware Filtering Method

B.1 Inputs and outputs

Inputs. (i) Base text T (raw OCR; base_revision=0); (ii) a set of correction events E in the schema of Table 6; and (iii) a trust policy π over event metadata.

Outputs. A provenance-filtered variant T_π constructed by replaying the subset of events selected by π , alongside an application trace (applied/skipped/conflicted events) for audit.

B.2 Policy language and policies used in this paper

A policy π is a predicate on events. We use: **All corrections** $\pi(e) \equiv \text{true}$; **Confidence threshold** $\pi_\tau(e) \equiv (\text{confidence}(e) \geq \tau)$ with $\tau \in \{0.50, 0.70, 0.85\}$ (Table 3) and $\tau = 0.70$ as the main provenance-filtered condition; and **Human-approved only** $\pi_{\text{approved}}(e) \equiv (\text{review_status}(e) = \text{approved})$.

We emphasize that confidence is treated as a within-pipeline ranking signal, not a calibrated probability.

B.3 Deterministic replay and overlap handling

Ordering. Given selected events $E_\pi = \{e \in E : \pi(e)\}$, we sort by span_start ascending, break

ties by event_id, and replay as base-anchored replacements.

Conflicts. Because all events are anchored to the same base revision, overlaps can be detected explicitly. When overlapping events are incompatible, we resolve deterministically under a fixed rule (e.g., prefer human over model over rule; prefer review_status=approved over unreviewed; otherwise defer to adjudication). The key methodological point is that overlap handling is *explicit and reportable*, rather than silently overwritten.