

The Patrologia Graeca Corpus: OCR, Annotation, and Open Release of Noisy Nineteenth-Century Polytonic Greek Editions

Chahan Vidal-Gorène^{1,2}, Bastien Kindt³

¹École nationale des chartes – PSL, Paris, France

²Calfa, Paris, France

chahan.vidal-gorene@chartes.psl.eu

³UCLouvain – CIOL Institut orientaliste, Louvain-la-Neuve, Belgium

bastien.kindt@uclouvain.be

Abstract

We present the *Patrologia Graeca Corpus*, the first large-scale open OCR and linguistic resource for nineteenth-century editions of Ancient Greek. The collection covers the remaining undigitized volumes of the *Patrologia Graeca* (PG), printed in complex bilingual (Greek–Latin) layouts and characterized by highly degraded polytonic Greek typography. Through a dedicated pipeline combining YOLO-based layout detection and CRNN-based text recognition, we achieve a character error rate (CER) of 1.05% and a word error rate (WER) of 4.69%, largely outperforming existing OCR systems for polytonic Greek. The resulting corpus contains around six million lemmatized and part-of-speech tagged tokens, aligned with full OCR and layout annotations. Beyond its philological value, this corpus establishes a new benchmark for OCR on noisy polytonic Greek and provides training material for future models, including LLMs.

Keywords: OCR, Ancient Greek, Patrologia Graeca, corpus release, lemmatization, polytonic Greek

1. Introduction

Despite major advances in digital philology, Ancient Greek scholarship still lacks a reliable, open, and linguistically enriched corpus covering the *Patrologia Graeca* (PG; 161 vols., 1st–15th c.). Compiled by Jacques-Paul Migne in Paris between 1857 and 1866, the PG reprints a vast range of patristic, exegetical, historical, hagiographical, legal, encyclopedic, poetic, and even narrative texts, encompassing much of Byzantine Greek literature. Yet some of these works have not been re-edited since the nineteenth century and remain unavailable in a machine-readable format. The series is widely accessible as PDF scans (e.g., patristica.net, archive.org, books.google.com portals or Roger Pearse’s blog) and lacks structural encoding, which severely limits textual reuse, searchability, and linguistic analysis (Migne, 1857-1866; Pearse, 2011).

The absence of a structured, open, and linguistically annotated corpus of the PG prevents the application of computational approaches to a significant portion of Greek literature, as well as the development of models representing the full diachronic and stylistic range of the language. While the [Perseus Digital Library](#) (Perseus 5; part of the Open Greek and Latin Project, with initial development led by Eldarion and funded by the Alexander von Humboldt Chair of Digital Humanities at Leipzig), the [Patristic Text Archive](#), and the [Thesaurus Linguae Graecae](#) have transformed access to ancient texts, they re-

main partial or restricted for the post-classical and Byzantine periods.

This project addresses these gaps by identifying, digitizing, and linguistically annotating the still-unavailable volumes of the PG. The resulting corpus—around 6,000,000 lemmatized and morphologically parsed words—substantially extends the lexical and stylistic range of available Ancient Greek resources. Thanks to the wide thematic scope and chronological span of the PG, the dataset contributes thousands of rare inflected forms, technical and theological terms, named entities, and toponyms that are underrepresented in existing corpora, thereby providing a stronger coverage for Ancient Greek language models in the future.

2. Previous Initiatives and the Challenges of the *Patrologia Graeca*

2.1. Digital initiatives around the PG

Previous attempts to digitize the *Patrologia Graeca* (PG) have primarily consisted in providing PDF scans without reliable text extraction or structural markup, notably through Migne (1857-1866) and related initiatives curated by Pearse. Beyond these, only partial textual coverage exists in scholarly databases such as the [Thesaurus Linguae Graecae](#) and in open-access repositories like the [Patristic Text Archive](#) (composed mainly of texts from the *Griechischen Christlichen Schriftsteller* series).

The most ambitious effort to date was the *Open Patrologia Graeca 1.0* (OPG) project (Crane et al., 2015). Launched in 2015 by Gregory Crane and collaborators within the Perseus Digital Library and the Open Greek and Latin consortium, it aimed to produce OCR-derived Greek and Latin texts for all 161 PG volumes—roughly 50 million words—using open-source engines such as OCRopus (Breuel, 2008) and Tesseract. The raw outputs, released openly on GitHub under OGL-PatrologiaGraecaDev Organization (2017), represented an unprecedented attempt at large-scale text recovery for late antique and Byzantine Greek. While the project demonstrated the feasibility of large-scale OCR for polytonic Greek, the output quality remained highly variable, with frequent segmentation errors (distinction of latin lines vs greek lines, but also segmentation or characters), inconsistent encoding, and minimal metadata. The data were never converted into TEI-structured or linguistically annotated form, and the repositories have remained unmaintained since 2017. Later analyses confirmed that the OCR accuracy on the PG corpus is typically below 90% and that recognition performance degrades sharply for small fonts and complex diacritics (Robertson and Boschetti, 2017; Varthis et al., 2022). These limitations illustrate both the promise and the fragility of fully automatic approaches to nineteenth-century Greek printing.

In summary, the digital landscape of the *Patrologia Graeca* remains fragmented: most volumes exist only as non-searchable PDF scans; large-scale OCR attempts have produced uncorrected and inconsistent text; and a few TEI-encoded editions cover isolated works.

2.2. OCR and Editorial Challenges

The *Patrologia Graeca* presents exceptional challenges for both OCR and subsequent linguistic analysis. Each volume reproduces heterogeneous nineteenth-century reprints, typically (but not constantly) arranged in two columns (Greek and Latin, often with overlapping content flowing from one language to the other), and surrounded by marginalia, running titles, and apparatus notes (see Fig. 1).

Typographical quality varies greatly from one volume to another, and the dense use of polytonic diacritics severely complicates line segmentation and character recognition (Varthis et al., 2022; Vidal-Gorène, 2023). Common scanning issues—page curvature, misalignment, and excessive binarization—further degrade image quality, increasing ambiguity in reading (see Fig. 2). Moreover, no standardized, page-aligned ground truth datasets exist for these materials, which prevents the training and evaluation of dedicated OCR or layout analysis models for complex polytonic Greek (Robertson and Boschetti, 2017; Romanello et al., 2021).

3. Related Work in Greek OCR and Document Analysis

Research on Greek OCR has primarily focused on the structural complexity of scholarly editions. Early work on the *Patrologia Graeca* (Robertson et al., 2014) addressed column segmentation through geometric preprocessing—detecting intercolumn regions via Hough transforms and filtering spurious letters with a k-nearest-neighbor classifier—thus stabilizing OCR input. Later studies on dense critical layouts (Sven and Matteo, 2022) confirmed that purely visual object-detection models (notably YOLO) outperform textual or hybrid approaches (LayoutLMv3, RoBERTa), since segmentation in such editions depends more on spatial than linguistic features.

Tesseract, trained mainly on synthetic data, performs poorly on the *Patrologia Graeca* (Crane et al., 2015; Varthis et al., 2022). Alternative approaches such as *word spotting* (Varthis et al., 2022) work well on specific cases but lack general-

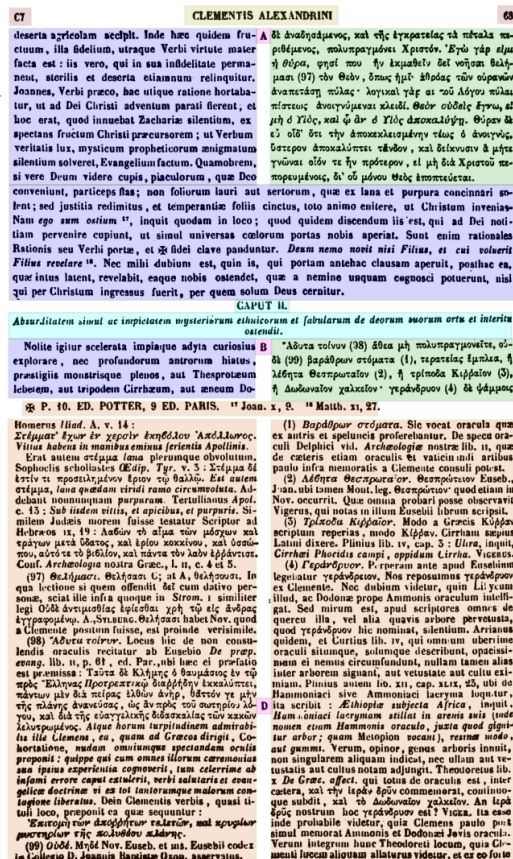


Figure 1: Example of page layout in PG with different semantic zones. Green regions correspond to the targeted Greek text.




Figure 2: Typography ambiguity in PG, variation of the character *alpha* with diacritics. See Section 5 for impact on OCR results

ization. More recently, page-level vision–language models (Semnani et al., 2025) have introduced contextual and multilingual OCR, capable of bypassing explicit segmentation.

Fine-tuning such models, however, requires extensive data, while existing corpora for polytonic Greek offer mostly line-level ground truth: *greek_cursive* from Von Stockhausen (2021); *gaza-iliad*, *voulgaris-aeneid*, and *gaza-batrachomyomachia* from Baumann (2022a,b, 2023); *GT4HistComment* (Romanello et al., 2021); *Pogretra* from Robertson (2021); and *GRPOLY-DB* (Gatos et al., 2015). These datasets provide strong baselines but remain narrowly focused (clean fonts, multilingual data; see Vidal-Gorène (2023)). Though unsuited to the thick and ambiguous typography of the *Patrologia Graeca*, *GT4HistComment* and *Pogretra* replicate more realistic scholarly layouts and supply useful lexical diversity for foundation model pretraining.

Automatic text recognition (OCR/HTR) is now standard in large-scale digitization, notably through *Transkribus* (Kahle et al., 2017), *eScriptorium*, and *Calfa Vision* (Vidal-Gorène et al., 2021). These platforms democratize model training by providing integrated fine-tuning pipelines based on baseline-oriented layout annotation and deep neural architectures, including fine-tunable CNN–LSTM and transformer-based networks (Reul et al., 2019; Vidal-Gorène et al., 2021). In particular, active-learning and automated iterative fine-tuning strategies implemented in *Calfa Vision* have yielded major gains on Greek, reducing CER by 63.94 points after only ten pages of transcription (Vidal-Gorène, 2023). The platform’s modular design also allows integration of YOLO-based layout models, identified as the most effective for these tasks (layout analysis and line detection, without baseline) (Sven and Matteo, 2022).

For lemmatization, POS tagging, and morphological parsing, we adopt the hybrid strategy developed for the Byzantine corpus *De Thessalonica Capta* (Kindt et al., 2022), combining neural tagging and rule- or dictionary-based post-correction. This method, built upon a fine-tuned version of

the PIE architecture (Manjavacas et al., 2019) adapted to the GREgORI tagset (Clérice, 2020), offers higher robustness for highly inflected and diachronic Greek. It reaches 97.94% lemmatization and 98.59% POS-tagging accuracy, while maintaining interpretability and stable handling of polylexical and enclitic forms.

This hybrid workflow also serves a corpus-building objective: it will yield approximately six million additional words of linguistically normalized Greek, thereby extending existing resources for model training. Such enrichment directly benefits the next generation of Ancient-Greek transformer models—e.g. *Ancient-Greek-BERT* (Singh et al., 2021), *GreBERTa* and *PhilBERTa* (Riemenschneider and Frank, 2023), or *LOGION* (Cowen-Breen et al., 2023)—whose performance depends on lexically diverse, diachronic data. Our corpus thus plays a dual role: improving automatic annotation quality through a dedicated tagset, and contributing new training material for future re-training or benchmarking of Ancient Greek BERT-like architectures.

Beyond OCR, recent work has explored the semantic enrichment of Greek and Latin corpora through translation alignment, syntactic annotation, and named entity recognition. Projects such as *Ugarit* and *Beyond Translation* (Crane et al., 2025) combine manual and automatic alignments with treebank data, enabling large-scale semantic and cross-lingual analysis. The creation of large, linguistically normalized corpora—such as the extended PG dataset proposed here—provides the substrate required for robust adaptation of these models to post-classical and Byzantine Greek.

4. Methodology and dataset for PG OCR and analysis

4.1. Data Preparation

To evaluate OCR performance on the *Patrologia Graeca* (PG), we first established a 30 pages test set, randomly sampled from the corpus and manually transcribed for the Greek text regions. Existing public models for Greek—namely the *tesseract-ocr*¹ and *Transkribus* model for nineteenth-century Greek prints²—fail to achieve an average character error rate (CER) below 5% (see Table 2). This confirms that PG presents challenges beyond the scope of existing models, due to its complex typography, heterogeneous layouts, and degraded scan quality.

¹https://github.com/tesseract-ocr/tessdata_best/blob/main/grc.traineddata

²<https://app.transkribus.org/models/public/text/19th-century-greek-8.0>

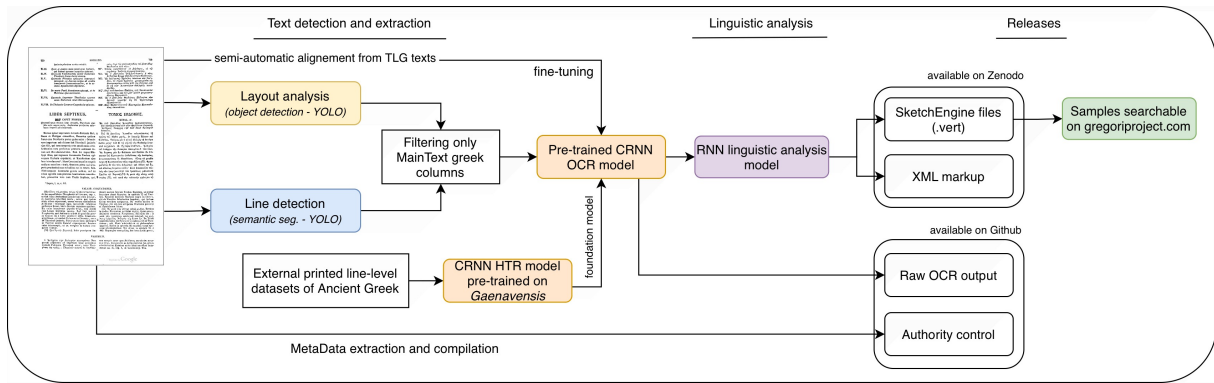


Figure 3: Overview of the OCR and annotation workflow for the *Patrologia Graeca* corpus.

4.2. Model Architecture and Fine-tuning

To address this, we used a pretrained word-based CRNN architecture on a Greek manuscript, the *Genavensis Græcus 44*. This model was then fine-tuned on printed datasets previously described, augmented with 50k synthetic samples generated through the *alumentations* library in order to create a foundation model for our task. We applied aggressive degradations (Gaussian noise, motion blur, compression artifacts, contrast shifts, elastic distortion, random fog, rain, and coarse dropout) to reproduce the visual variability and damage patterns typical of PG scans. While this foundation model did not yet reduce the CER significantly, it stabilized the word error rate (WER) around 10% thanks to improved lexical coverage, and served as a robust starting point for PG-specific fine-tuning (see Table 2).

The main PG training dataset was then created following the iterative fine-tuning strategy of Vidal-Gorène (2023). Segments of the PG already available in digital form through the *Thesaurus Linguae Graecae* were semi-automatically aligned within the Calfa Vision platform (Vidal-Gorène et al., 2021), combining automatic layout detection with manual correction. The resulting dataset comprises 445 images representative of the typographical and structural diversity of the PG (see Table 1). A specific focus has been made on pages with crossing text-regions (see Figure 1 with the Latin column in blue).

For text analysis, we use a model trained on a representative sample of the corpora from the GREgORI database (Kindt, 2025), which comprises more than two million lemmatized word forms covering Classical and Byzantine Greek from the first to the fifteenth century, lemmatized and morphosyntactically annotated through rule-based dictionary methods and manually corrected. The model used is the one trained by Kindt et al. (2022) for the *De Thessalonica Capta* corpus. Figure 3 illustrates the processing pipeline.

5. Results and Corpus Release

5.1. OCR and Layout Performance

The results in Table 2 summarize the performances in text recognition and layout analysis. Our OCR model, specifically fine-tuned for the PG material, clearly outperforms all previously used systems on this corpus, as well as a model trained on synthetic data, achieving a gain of 5–7 points in CER and 6–10 points in WER.

The most frequent OCR errors concern diacritic confusions among polytonic Greek characters. The majority of mismatches [e.g., $\acute{\iota}$ (03Af) / $\dot{\iota}$ (1F77), $\acute{\alpha}$ (03AC) / $\dot{\alpha}$ (1F71), $\acute{\epsilon}$ (03AD) / $\dot{\epsilon}$ (1F73), $\acute{\omicron}$ (03CC) / $\dot{\omicron}$ (1F79), $\acute{\upsilon}$ (03CD) / $\dot{\upsilon}$ (1F7B), $\acute{\eta}$ (03AE) / $\dot{\eta}$ (1F75), $\acute{\omega}$ (03CE) / $\dot{\omega}$ (1F7D)] correspond to visually identical characters whose accent encodings differ, typically between monotonic and polytonic forms. This is due to non-homogeneity in the datasets used. These substitutions account for over 80% of all errors, showing that the model usually recognizes the correct letter identity but fails to reproduce the exact diacritic.

A closer examination of the 123 distinct error pat-

Global information	Count
Images	445
TextLines	11,096
LinePolygons	6,921
TextRegions	4,342
Semantic categories	Count
MainText_ColGreek	1,022
MainText_ColLatin	761
MainText_Title	267
Marginalia	476
Marginalia_Footnote	652
Marginalia_PageNumber	611
Marginalia_ParagraphNumber	248
Title_RunningTitle	305

Table 1: OCR ground-truth

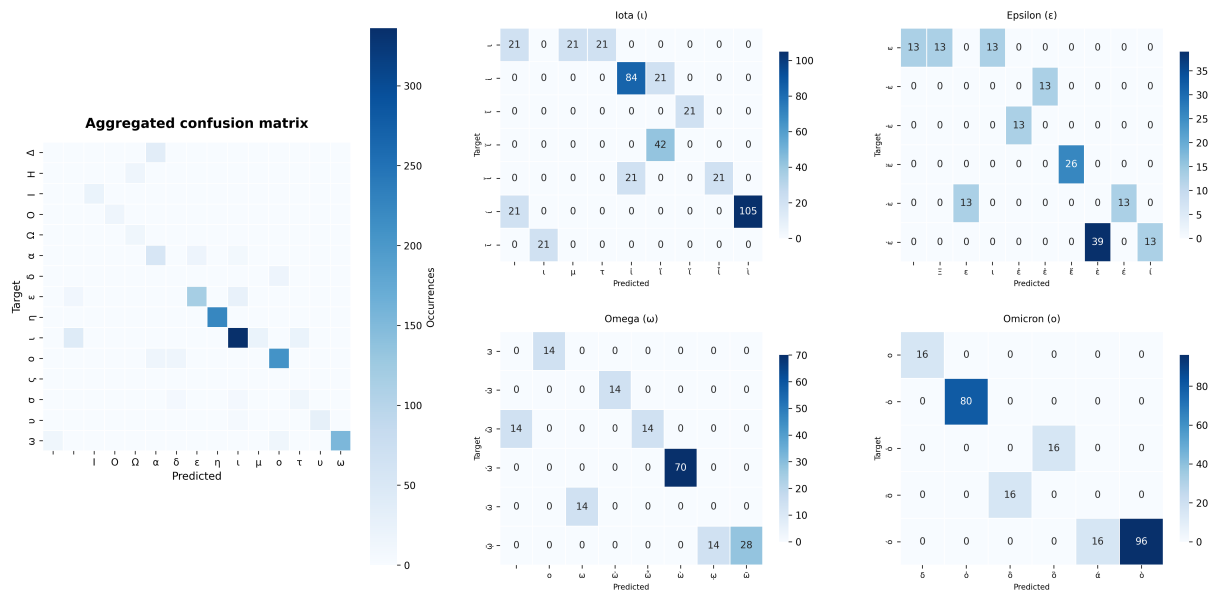


Figure 4: Character-level confusion matrix (left) and detailed confusion patterns for iota, omega, epsilon, and omicron showing diacritic variation

terms observed across 1,572 total occurrences reveals systematic challenges by letter group (Fig. 4). *Iota* exhibits the highest error rate (315 occurrences), primarily due to confusion between rough and smooth breathing marks (e.g., different diacritics on ι), followed by omission of diacritics entirely. Interestingly, *iota* is also occasionally misread as morphologically unrelated characters such as τ or μ , suggesting low-level visual similarity in poorly inked regions. *Omicron* (144 occurrences) shows predominantly breathing-related errors, with a few instances of cross-letter confusion with δ or misrecognition of uppercase forms. *Omega* (126 occurrences) presents similar diacritic instability, especially in subscript *iota* and accent type, with one notable cross-vowel error reflecting visual proximity in compressed typography. *Epsilon* (117 occurrences) exhibits the most diverse error patterns: breathing confusions, total diacritic loss, and a striking confusion with uppercase Ξ likely due to similar vertical strokes in certain typefaces.

A smaller set of errors involves spacing and punctuation, final sigma versus punctuation ($\varsigma \rightarrow$ comma or period), inter-consonant confusions (σ/δ , σ/τ , ρ/o), and uppercase mismatches (Δ/Λ , Ψ/T), accounting for fewer than 10% of total errors but remaining problematic for lemmatization pipelines.

Overall, most transcription errors can be mitigated through systematic Unicode normalization before training and post-correction rules. This leads to a substantially lower CER. Errors are still more frequent in uppercase characters (therefore mainly titles), which would require specific reinforcement during training. Post-processing is a necessary task, because if the symbol appears similar, it

penalizes text analysis and keyword research. Only Greek columns and titles are relevant in our case; other regions are excluded from the pipeline.

Text recognition

Model	CER (%)	WER (%)
Tesseract (Greek)	11.57	39.65
Transkribus (19th c. Greek)	6.14	14.82
Pre-trained CRNN (with artificial noise)	8.12	11.23
Ours (PG fine-tuned)	1.05	4.69

Layout detection

Class	P	R	mAP50
MainText_ColGreek	0.963	0.994	0.973
MainText_ColLatin	0.969	1.000	0.982
MainText_Title	0.462	0.950	0.462
Marginalia	0.422	0.891	0.427
Marginalia_Footnote	0.493	0.975	0.511
Marginalia_PageNumber	0.382	0.783	0.380
Marginalia_ParagraphNumber	0.353	0.679	0.408
Running	0.477	0.883	0.542

Line detection and reading order

Task	P	R	mAP50
Line detection	0.983	0.994	0.973
Reading order	0.98	-	-

Table 2: OCR performance (layout and recognition) comparison on the 30-page PG test set.

Beyond the quantitative results on the test set, applying the models in real production conditions reveals additional sources of error. Although it occurs only occasionally, a few cross-column regions in Latin may be mistakenly detected, introducing noticeable noise in the subsequent text recognition. The most persistent challenge, however, concerns title areas, which are not always clearly distinguished between Greek and Latin

scripts. Another issue sometimes observed in line detection—particularly in overlapping or crossing regions—is the splitting of a single line into two fragments with their reading order reversed. To mitigate these residual problems, we adopt an active learning strategy, continuously correcting and reinjecting misrecognized samples into the training data to progressively improve model robustness.

After applying the OCR models to the entire set of PG volumes and processing the recognized text, we obtain a corpus of approximately six million words.

5.2. Corpus Characteristics and Visualization

To evaluate both the linguistic and visual characteristics of the corpus we release, we projected the existing open corpora into two embedding spaces (Fig. 5). Subfigure (a) shows a t-SNE visualization of sentence-level embeddings derived from *fastText* models trained on Ancient Greek. The PG corpus (yellow) forms a coherent but distinct cluster, confirming its lexical specificity compared to other datasets. In contrast, the Pogretra data (grey) appear more dispersed, reflecting OCR noise and inconsistent orthography, while *Gaza-Batrachomyomachia* and *Greek_cursive* occupy peripheral, stylistically homogeneous regions.

Subfigure (b) presents a t-SNE projection of line-level image embeddings extracted from VGG16 (pretrained CNN). Unlike in the lexical space, the visual distribution shows strong inter-corpus separation, highlighting typographic and palaeographic variation. The PG lines occupy a dense and compact region, visually distinct from both manuscript and modern print corpora. These observations underline the atypical nature of the PG—typographically complex, lexically rich, and divergent from available Greek OCR datasets—thus motivating the creation of a dedicated resource and the development of tailored OCR, normalization workflows and motivating the release of this new corpus for Ancient Greek studies.


5.3. Output Format and Public Release

Prior to linguistic analysis, OCR-processed texts undergo a preparation step involving the removal of hyphens, empty lines, and Latin characters, words, or lines not belonging to the target Greek text. As it stands, however, this operation is still carried out manually, without any assurance of completion. Following this preparation, the output is structured in the `.vert` format used by Sketch Engine, where each word token (`<w>`) is annotated with five layers of information: the OCR wordform, the lowercase diacritic-free form (intuitive form), the lemma, the

lowercase diacritic-free lemma (intuitive lemma), and a morpho-syntactic tag (see Fig. 6). The intuitive form and intuitive lemma allow users to submit lexical queries using plain lowercase characters, free from diacritical marks or accents—a practical affordance for researchers less familiar with polytonic Greek orthography.

The XML structure proposed here could naturally be adapted to other environments and use cases as needed. Beyond annotation, the format preserves document, page, and line identifiers, ensuring full traceability between OCR output and its original textual context—an essential feature for philological and linguistic exploration of the PG corpus. This hierarchical representation enables direct import into corpus analysis tools such as Sketch Engine or TreeTagger-compatible environments.

The resulting corpus is directly searchable on gregoriproject.com, while the raw OCR data and structured `.vert` files are publicly released on Github to ensure transparency and reproducibility. This release includes the full OCR text, line segmentation, and layout annotations, facilitating



(a) t-SNE of *fastText* embeddings (semantic)



(b) t-SNE on VGG16 embeddings (vision)

Figure 5: Comparison of semantic and visual distributions of ancient Greek corpora.

further experimentation and benchmarking on polytonic Greek material.

```

1 <doc id="PG 134 - AI processed">
2 <w id="PDF Page 31, Line 1 - word id 1">
3 ΠΡΟΟΙΜΙΟΝ.προοιμιον προοίμιον προοιμιον N+Com
4 </w>
5 <w id="PDF Page 31, Line 2 - word id 2">
6 Α'. α ἄλφα αλφα N+Lettre
7 </w>
8 <w id="PDF Page 31, Line 2 - word id 3">
9 Εύστόχως ευστοχως εύστόχως ευστοχως I+Adv
10 </w>
11 <w id="PDF Page 31, Line 2 - word id 4">
12 ἄν αν ἄν αν I+Part
13 </w>
14 <w id="PDF Page 31, Line 2 - word id 5">
15 τις τις τις τις PRO+Ind
16 </w>
17 <w id="PDF Page 31, Line 2 - word id 6">
18 εἴποι, ειποι λέγω λεγω V
19 </w>
20 <w id="PDF Page 31, Line 2 - word id 7">
21 ἐπιτωθάζων επιτωθαζων ἐπιτωθάζω επιτωθαζωV
22 </w>
23 <w id="PDF Page 31, Line 2 - word id 8">
24 μοι, μοι ἐγώ εγω PRO+Per1s

```

Figure 6: Example of the structured output in .vert format, with OCR wordform, intuitive form (clean), lemma, intuitive lemma (clean), and POS, with line IDs, word IDs and reference to PDF page.

6. Conclusion

This study presented an OCR, text normalization, and text analysis workflow dedicated to the still undigitized texts of the *Patrologia Graeca*, addressing the specific challenges of poorly printed nineteenth-century polytonic Greek and complex bilingual layouts. The proposed pipeline is lightweight by design, combining YOLO-based layout detection with iterative CRNN fine-tuning on a modest PG-specific ground truth, yet achieves a CER of 1.05% and a WER of 4.69%, improving by more than five points over existing CER baselines. In addition to OCR and layout annotations, the project releases a six-million-word lemmatized corpus, openly available for reuse (silver corpus). Its value lies not only in its scale but in its lexical and stylistic distinctiveness. These resources provide a consistent and verifiable basis for further work in historical OCR, text normalization, and the adaptation of large language models to Ancient Greek.

7. Data Availability

Annotations for OCR training and the corpus (raw outputs, and VERT files) are available on <https://github.com/calfa-co/Patrologia-Graeca>. Currently, only sample VERT files are provided due to the large size of the repository. The full set is released on Zenodo (<https://zenodo.org/records/15780625>). The corpus is freely accessible on <https://www.gregorioproject.com>.

8. Acknowledgements

This research was supported by: ASBL Byzantion, Calfa (Paris), UCLouvain - CIOL - Centre d'études orientales - Institut orientaliste de Louvain, UCLouvain - FSS - Fondation Sedes Sapientiae, UCLouvain GREgORI Project, UCLouvain - INCAL - Institut des Civilisations Arts et Lettres UCLouvain - RSCS - Institut de recherche pluridisciplinaire Religions Spiritualités Cultures Sociétés, and other (anonymous) private funders. We thank Professor Jean-Marie Auwers (UCLouvain, RSCS), Professor Sébastien Moureau (UCLouvain/CIOL) and Doctor Véronique Somers (UCLouvain/CIOL) for their contributions to data curation, annotation, and evaluation.

9. Bibliographical References

- Thomas M Breuel. 2008. The ocropus open source ocr system. In *Document recognition and retrieval XV*, volume 6815, pages 120–134. SPIE.
- C. Cowen-Breen, C. Brooks, J. Haubold, and B. Graziosi. 2023. [Logion: Machine learning for greek philology](#).
- Gregory Crane, Alison Babeu, and Farnoosh Shamsian. 2025. [Greek, latin and augmented intelligence: The other ai](#). *The Classical Review*, 75(1):7–22.
- Gregory Crane, Federico Boschetti, Matt Munson, Bruce Robertson, and Nick White. 2015. Open Patrologia Graeca 1.0. <https://sites.tufts.edu/perseusupdates/2015/08/07/open-patrologia-graeca-1-0/>. Blog announcement within the Open Greek and Latin Project.
- Basilis Gatos, Nikolaos Stamatopoulos, Georgios Louloudis, Giorgos Sfikas, George Retsinas, Vassilis Papavassiliou, Fotini Sunistira, and Vassilis Katsouras. 2015. [Gpoly-db: An old greek polytonic document image database](#). In *2015*

- 13th International Conference on Document Analysis and Recognition (ICDAR), pages 646–650.
- Philip Kahle, Sebastian Colutto, Günter Hackl, and Günter Mühlberger. 2017. Transkribus—a service platform for transcription, recognition and retrieval of historical documents. In *2017 14th iapr international conference on document analysis and recognition (icdar)*, volume 4, pages 19–24. IEEE.
- Bastien Kindt. 2025. [The Online Corpus of the GREgORI Project: Byzantine Greek Corpus – Corpus Description](#). Unpublished internal report. Version 2025-01. © UCLouvain – GREgORI Project 2024.
- Bastien Kindt, Chahan Vidal-Gorène, and Saulo Delle Donne. 2022. Analyse automatique du grec ancien par réseau de neurones. évaluation sur le corpus de thessalonica capta. *Bulletin de l'Académie Belge pour l'Étude des Langues Anciennes et Orientales*, pages 537–562.
- Enrique Manjavacas, Ákos Kádár, and Mike Kestemont. 2019. Improving lemmatization of non-standard languages with joint learning. *arXiv preprint arXiv:1903.06939*.
- Christian Reul, Dennis Christ, Alexander Hartelt, Nico Balbach, Maximilian Wehner, Uwe Springmann, Christoph Wick, Christine Grundig, Andreas Büttner, and Frank Puppe. 2019. [Ocr4all—an open-source tool providing a \(semi-\)automatic ocr workflow for historical printings](#). *Applied Sciences*, 9(22):4853.
- Frederick Riemenschneider and Anette Frank. 2023. [Exploring large language models for classical philology](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL'23)*. Association for Computational Linguistics, Toronto, Canada. To appear.
- Bruce Robertson and Federico Boschetti. 2017. Large-scale optical character recognition of ancient greek. *Mouseion*, 14(3):341–359.
- Bruce Robertson, Christoph Dalitz, and Fabian Schmitt. 2014. Automated page layout simplification of patrologia graeca. In *Proceedings of the First International Conference on Digital Access to Textual Cultural Heritage*, pages 167–172.
- Matteo Romanello, Sven Najem-Meyer, and Bruce Robertson. 2021. [Optical character recognition of 19th-century classical commentaries: The current state of affairs](#). In *Proceedings of the ACM Symposium on Document Engineering (DocEng '21)*. ACM.
- Sina J Semnani, Han Zhang, Xinyan He, Merve Tekgürler, and Monica S Lam. 2025. Churro: Making history readable with an open-weight large vision-language model for high-accuracy, low-cost historical text recognition. *arXiv preprint arXiv:2509.19768*.
- Pranaydeep Singh, Gorik Rutten, and Els Lefever. 2021. A pilot study for bert language modelling and morphological analysis for ancient and medieval greek. In *The 5th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature (LaTeCH-CLfL 2021)*.
- Najem-Meyer Sven and Romanello Matteo. 2022. Page layout analysis of text-heavy historical documents: a comparison of textual and visual approaches. *arXiv preprint arXiv:2212.13924*.
- Evangelos Varthis, Spyros Tzanavaris, Ilias Giarenis, Sozon Papavlasopoulos, Manolis Drakakis, and Marios Poulos. 2022. Semantic enrichment on large scanned collections through their “satellite texts”: the paradigm of migne’s patrologia graeca. *Information Discovery and Delivery*, 50(2):217–234.
- Chahan Vidal-Gorène. 2023. La reconnaissance automatique d’écriture à l’épreuve des langues peu dotées. *The Programming Historian en français*, (5).
- Chahan Vidal-Gorène, Boris Dupin, Aliénor Decours-Perez, and Thomas Riccioli. 2021. A modular and automated annotation platform for handwritings: evaluation on under-resourced languages. In *International Conference on Document Analysis and Recognition*, pages 507–522. Springer.

10. Language Resource References

- Ryan Baumann. 2022a. Gaza-iliad. <https://github.com/ryanfb/kraken-gaza-iliad>. Original edition 2019. GitHub repository.
- Ryan Baumann. 2022b. Voulgaris-aeneid. <https://github.com/ryanfb/kraken-voulgaris-aeneid>. Original edition 2019. GitHub repository.
- Ryan Baumann. 2023. Gaza-batrachomyomachia. <https://github.com/ryanfb/kraken-gaza-batrachomyomachia>. Original edition 2020. GitHub repository.

Thibault Clérice. 2020. Pie Extended, an extension for Pie with pre-processing and post-processing. <https://doi.org/10.5281/zenodo.3883589>.

Jacques P. Migne. 1857-1866. Patrologiae Curcus Completus. Series Graeca (PG) PDFs. <https://patristica.net/graeca/>. Accessed 2025-10-05.

OGL-PatrologiaGraecaDev Organization. 2017. Patrologia graeca ocr repositories. <https://github.com/OGL-PatrologiaGraecaDev>. Accessed 2025-10-05.

Patristic Text Archive. Patristic text archive (pta). <https://pta.bbaw.de/en>. Accessed 2025-10-05.

Roger Pearse. 2011. Patrologia graeca (pg) pdfs. <https://www.roger-pearse.com/weblog/patrologia-graeca-pg-pdfs/>. Accessed 2025-10-05.

Perseus Digital Library. Scaife Viewer. <https://scaife.perseus.org/>. Accessed 2025-10-06.

Bruce Robertson. 2021. Pogretra: Polytonic greek training data from historic texts.

Thesaurus Linguae Graecae. The abridged tlg. <https://stephanus.tlg.uci.edu>. Accessed 2025-10-05.

Annette Von Stockhausen. 2021. ocr-greek_cursive. https://github.com/pharos-alexandria/ocr-greek_cursive.

Appendices

The PG comprises 161 volumes, numbered from 1 to 161 (volumes 7, 16, 86 and 87 being divided into two or three parts). All these volumes are available in PDF format (files are listed, e.g., on the portal Patristica.net or on the blog roger-pearse.com). At the beginning of the project, 55 volumes were selected, based on their presence, or not, in the TLG corpus. As of now, some texts have been excluded because they were in the meantime uploaded into the TLG corpus, or because their layout was quite specific, with more than two columns or including languages other than Latin or Greek, such as Syriac (see, e.g., vol. 1; <https://books.google.be/books?id=qxANpCDQIjIC>, p. 102 and 198). The 32 volumes that were ultimately selected and are currently being processed in the project are listed in Table 3 below.

PG Vol.	URL	Date	Word Count
3	books.google.com/.../ALFUAAAAMAAJ	5th-6th AD (?)	134,866
5	books.google.com/.../PIQe9iqBoeQC	1st-2d AD	46,164
6	books.google.com/.../NZLYAAAAMAAJ	2d AD	170,482
8	books.google.com/.../BwcRAAAAYAAJ	2d-3rd AD	168,277
9	books.google.com/.../QAgRAAAAYAAJ	2d-3rd AD	82,135
16.3	books.google.com/.../mAsRAAAAYAAJ	3rd AD	60,921
21	books.google.com/.../BwcRAAAAYAAJ	3rd-4th AD	236,625
42	books.google.com/.../QHYRAAAAMAAJ	4th-5th AD	161,237
67	books.google.com/.../WivbHxo0L-sC	4th-6th AD	170,445
71	books.google.be/.../worYAAAAMAAJ	4th-5th AD	210,957
73	books.google.com/.../ywsNQz1fTewC	4th-5th AD	191,303
87.1	books.google.com/.../CmcUAAAQAQAJ	5th-6th AD	151,167
107	books.google.com/.../Ru4GAAAQAQAJ	9th-10th AD	196,727
109	books.google.com/.../Z0naYVT0w-EC	9th-10th AD	148,584
112	books.google.com/.../nyNKAAAACAAJ	10th AD	129,556
113	books.google.com/.../Z_UAAAQAQAJ	10th AD	104,371
118	books.google.com/.../_j_UAAAQAQAJ	11th AD (?)	208,448
121	books.google.com/.../PIQe9iqBoeQC	11th AD	160,853
122	archive.org/.../patrologiaecurs61migngoog	11th AD	150,647
123	books.google.com/.../-SFJAAAACAAJ	11th-12th AD	208,024
124	books.google.com/.../AccUAAAQAQAJ	11th-12th AD	210,302
125	books.google.com/.../Z7_UAAAAMAAJ	11th-12th AD	172,696
126	books.google.com/.../eTYRAAAAYAAJ	11th-12th AD	164,706
134	books.google.com/.../DrvUAAAAMAAJ	11th-12th AD	196,859
139	books.google.com/.../lscUAAAQAQAJ	12th-13th AD	134,703
146	books.google.com/.../xCKAAAACAAJ	13th-14th AD	156,848
148	books.google.com/.../l1Ww11vmdb4C	13th-14th AD	234,855
151	books.google.com/.../PyNKAAAACAAJ	13th-14th AD	399,518
153	books.google.com/.../dIYAAAAMAAJ	13th-14th AD	230,239
155	books.google.com/.../_McUAAAQAQAJ	14th-15th AD	175,482
157	books.google.com/.../YB8RAAAAYAAJ	15th AD	95,020
158	books.google.com/.../hBIT951CqEsC	12th AD	163,148
Total			5,605,015

Table 3: List of processed volumes of the PG, url of the PDF files, date of the texts, and wordcount.