

RESEARCH

Open Access



From chaos to clarity: schema-constrained AI for auditable biomedical evidence extraction from full-text PDFs

Pouria Mortezaagha^{1,3*}, Joseph Shaw², Bowen Sun³ and Arya Rahgozar^{1,3}

Abstract

Background Biomedical evidence synthesis depends on accurate extraction of methodological, laboratory, and outcome variables from full-text research articles. These variables are predominantly embedded in complex scientific PDFs that interleave multi-column text, tables, figures, and captions, making manual abstraction time-intensive, error-prone, and increasingly impractical at the scale of contemporary systematic reviews. Despite advances in layout-aware and multimodal document models, end-to-end extraction systems suitable for evidence synthesis remain constrained by limited throughput, OCR error propagation, and insufficient audibility.

Methods We propose a schema-constrained AI extraction system that transforms full-text biomedical PDFs into structured, analysis-ready records by explicitly restricting model inference through typed schemas, controlled vocabularies, and evidence-gated decisions. Documents are ingested using resume-aware hashing, partitioned into page-level and caption-aware chunks, and processed asynchronously under explicit concurrency and rate-limiting controls. A high-accuracy OCR model is guided by multiple domain-specific schemas covering bibliographic metadata, study design, populations, laboratory assays, timing and thresholds, clinical outcomes, and diagnostic performance. Chunk-level outputs are deterministically merged into study-level records using controlled vocabularies, conflict-aware handling of scalar fields, set-based aggregation of list-valued fields, and sentence-level evidence capture to enable traceability and post-hoc audit.

Results Applied to a corpus of 734 biomedical articles on direct oral anticoagulant (DOAC) level measurement, the pipeline processed all documents without manual intervention while maintaining stable throughput. Schema-constrained extraction exhibited strong internal consistency, with sentence-level provenance populated for nearly all supported decisions. Iterative schema and prompt refinement yielded substantial improvements in extraction fidelity, particularly for outcome definitions, assay classification, and global coagulation testing. Outputs included reproducible CSV/Parquet datasets and caption-aware multimodal markdown reconstructions supporting efficient expert review.

Conclusions Schema-constrained AI extraction enables scalable and auditable extraction of structured evidence from heterogeneous scientific PDFs. By combining deterministic chunking, asynchronous orchestration, controlled vocabularies, sentence-level provenance, and aggregated analytical outputs, the proposed pipeline aligns modern

*Correspondence:
Pouria Mortezaagha
pmortezaagha@ohri.ca

Full list of author information is available at the end of the article



© The Author(s) 2026. **Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

document understanding capabilities with the transparency, reproducibility, and reliability demands of biomedical evidence synthesis.

Keywords Biomedical evidence synthesis, Scientific PDF information extraction, OCR-based document understanding, Schema-constrained AI, Provenance-aware extraction, Auditability, Systematic reviews, Biomedical natural language processing, Multimodal document analysis

Introduction

Biomedical evidence synthesis increasingly relies on transforming large volumes of full-text scientific literature into structured, analysis-ready representations. While abstracts and curated databases provide partial coverage, essential methodological and clinical details such as study design, patient populations, laboratory assays, decision thresholds, and outcome definitions are reported primarily within the full text and are dispersed across narrative sections, tables, figures, and captions embedded in complex document layouts (Fig. 1), rendering manual abstraction labor-intensive, error-prone, and fundamentally unscalable.

This challenge is no longer marginal. The volume of biomedical publications now exceeds what can be realistically processed by human reviewers alone, even within narrowly defined clinical domains [1]. Systematic reviews and guideline updates increasingly confront thousands of full-text articles, each requiring careful interpretation of heterogeneous reporting practices. As a result, evidence synthesis workflows remain bottlenecked by manual annotation [2], despite substantial progress in document understanding and natural language processing, and a growing ecosystem of tools designed to improve systematic review efficiency [3].

Automating this process is complicated by the dominance of the PDF format, which preserves visual appearance rather than semantic structure. Scientific PDFs encode text as positioned glyphs rather than logical units, and routinely employ multi-column layouts, floating tables and figures, dense numerical reporting, and caption-embedded methodological information. Errors in reading-order reconstruction, table detection, or caption association can cascade into downstream extraction failures, particularly in biomedical contexts where small numerical or categorical inaccuracies materially affect interpretation. Reliable extraction from scientific PDFs, therefore, requires not only accurate text recognition but robust layout awareness, semantic grounding, and mechanisms to control inference under ambiguity. Figure 2 summarizes the structural and layout heterogeneity of scientific PDFs that systematically hinders scalable and auditable evidence extraction.

Recent advances in layout-aware transformers, multimodal document models, and neural transcription systems have substantially improved the structural and semantic parsing of scientific documents. However, these

models remain sensitive to OCR noise, layout heterogeneity, and domain shift, and often lack explicit guarantees of traceability. In high-stakes biomedical applications, accuracy alone is insufficient: extracted variables must be auditable, reproducible, and explicitly linked to supporting evidence in the source document.

These challenges are especially acute in laboratory-based clinical research, where reporting heterogeneity is high and key variables are frequently expressed in non-narrative form. Studies of direct oral anticoagulant (DOAC) level measurement exemplify this complexity, reporting diverse assay families, pre-analytical conditions, sampling times, diagnostic thresholds, and outcome definitions that are often embedded in tables or figure captions. Scalable and standardized extraction of such information is essential for comparative evaluation and evidence synthesis, yet remains impractical with manual workflows given current publication volumes [4].

This work addresses the following research question: *Can AI systems be designed to extract structured evidence from full-text biomedical PDFs at scale while explicitly constraining inference to ensure auditability across heterogeneous document layouts?*

Specifically, we aim to achieve scale without sacrificing controllability by constraining outputs to typed schemas and attaching sentence-level provenance to key variables.

To this end, we propose an end-to-end, OCR-driven pipeline that couples high-accuracy document transcription using Mistral OCR 3 (`mistral-ocr-2512`), a multimodal vision-language model, with schema-constrained structured extraction. Mistral OCR 3 was selected for its native support for schema-constrained structured output via JSON schema enforcement, high-accuracy multimodal page understanding with layout awareness, and competitive throughput under API-based deployment. Documents are processed at the page level using deterministic chunking and asynchronous orchestration, enabling stable large-scale execution under service constraints. Domain-specific schemas restrict outputs to typed fields with controlled vocabularies and explicit negative rules, while sentence-level evidence is captured for key decisions to ensure traceability and post-hoc auditing.

We evaluate the pipeline on a corpus of biomedical research articles related to DOAC level measurement, spanning randomized trials, observational studies, diagnostic accuracy studies, and systematic reviews. The results demonstrate reliable corpus-scale processing,

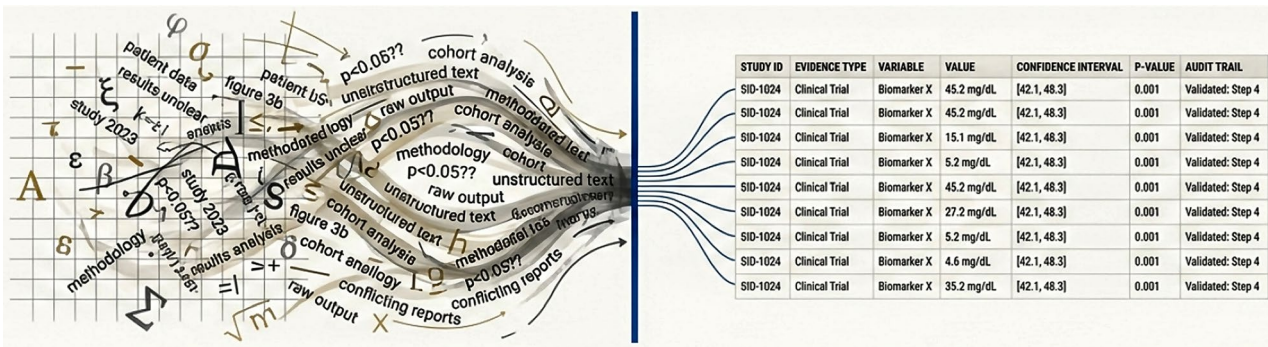


Fig. 1 From unstructured scientific PDFs to structured, provenance-linked evidence

Fig. 2 Structural challenges of scientific PDFs for evidence extraction

high internal consistency across extraction payloads, and substantial improvements in practical extraction fidelity following iterative schema and prompt refinement. By integrating layout-aware OCR, explicit schema constraints, and evidence-centric design, this work presents a practical and extensible framework for automated biomedical literature annotation that is aligned with the scale, rigor, and transparency demands of modern evidence synthesis.

ContributionsThe contribution of this work is not any single component in isolation, but the principled integration of schema constraints, evidence gating, and deterministic merging into a reproducible end-to-end workflow that addresses the specific failure modes of biomedical evidence synthesis. Specifically:

- We propose a *schema-constrained AI extraction system* that governs long-document inference through typed schemas, controlled vocabularies, and explicit negative rules, reducing over-inference under heterogeneous layouts and OCR noise.
- We formalize a *deterministic chunking and consolidation strategy* that reconciles page-level

annotations into coherent study-level records, enforcing consistency for scalar fields, set-based aggregation for list-valued fields, and explicit conflict detection.

- We treat *sentence-level provenance as a first-class output*, binding synthesis-critical variables to explicit supporting text to enable transparent audit, expert adjudication, and reproducible refinement.
- We validate the approach through corpus-scale evaluation in a challenging laboratory medicine domain (DOAC level measurement), demonstrating stable throughput under service constraints and improved extraction fidelity via expert-in-the-loop schema refinement.

Related work

Scientific knowledge remains predominantly disseminated as PDF, a format optimized for visual presentation rather than machine-readable structure. Information extraction (IE) from scientific PDFs is inherently a multi-stage problem requiring reconciliation of visual layout, symbolic text, and document semantics. Progress has

followed a trajectory from heuristic layout analysis to neural approaches that integrate layout with language.

Structural reconstruction and early scholarly PDF extraction

Most scholarly PDF IE pipelines begin with *structural reconstruction*, recovering spatial text units before segmenting content. Systems such as GROBID, CERMINÉ, and ParsCit use layout cues and learned models to extract bibliographic metadata and references [5–7], but rely on handcrafted features that degrade under heterogeneous layouts. A central design choice is whether to treat PDFs as *text-first* (preserving exact text from born-digital PDFs) or *image-first* (processing via OCR). These tradeoffs motivate hybrid approaches combining OCR with explicit constraints.

From heuristic rules to feature-based learning

Early scholarly PDF IE relied on heuristic layout analysis using rules over font size, indentation, and keyword patterns [8, 9], but such systems degrade under distribution shift. Feature-based machine learning improved generalizability: GROBID and ParsCit employ CRFs with engineered features [5, 7, 10], and CERMINÉ uses a modular workflow for metadata extraction [6]. However, these approaches still depend on task-specific feature design and struggle with long-range dependencies. Large-scale corpora such as S2ORC highlighted the importance of reliable PDF-to-structure conversion for scientific text mining [11].

Deep learning for layout and multimodal structure

Deep learning reduced reliance on manual feature engineering for document layout analysis. Systems such as PDFFigures 2.0 target figure, table, and caption extraction [12, 13], while VILA leverages visual layout groupings for structured extraction [14]. Datasets such as PubLayNet, DocBank, and OmniDocBench provided supervision for training scientific layout models [15–17].

Layout-aware transformers and relational document modeling

A major advance in document understanding has been the integration of spatial layout into transformer-based language models. LayoutLM introduced joint modeling of token text and two-dimensional position embeddings [18], and LayoutLMv2 extended this paradigm by incorporating visual features to improve cross-modal alignment in visually rich documents [19]. DocFormer further unified textual, spatial, and visual representations within a single end-to-end architecture [20].

Subsequent work expanded toward *layout-aware generative modeling*. DocLLM conditions generative language models on bounding-box embeddings to produce structured text that respects document layout without heavy visual encoders [21], while GraphDoc models documents as relational graphs to explicitly encode reading order, hierarchy, and dependencies between regions

[22]. These approaches reduce structural ambiguity arising from multi-column layouts and long-range dependencies, but they remain sensitive to domain shifts, scanned document quality, and heterogeneous formatting conventions.

Semantic extraction with scientific language models

Layout reconstruction alone is insufficient for biomedical evidence synthesis, which requires reliable extraction of entities and relations. Pretrained scientific language models such as SciBERT improve this semantic layer [23–25], but document pipelines often decouple layout parsing from semantic extraction, and errors in either stage can undermine downstream reliability.

OCR-driven and end-to-end transcription approaches

Recent work has revisited OCR itself as a learning problem for scientific documents containing complex typography, symbols, and mathematics. Nougat proposes a visual transformer that converts scientific PDFs into structured markup-like representations, aiming to recover semantic structure lost in traditional OCR pipelines [26]. Related models such as TextMonkey explore end-to-end neural transcription through vision–language pretraining [27], while subsequent work improves long-document fidelity and layout robustness for scientific OCR [28]. More recently, olmOCR demonstrates large-scale PDF-to-text conversion using vision-language models, achieving high throughput on diverse document collections [29]. These approaches blur the boundary between transcription and semantic understanding.

OCR-free document understanding models, such as Donut, bypass explicit OCR by directly generating structured outputs from document images [30]. While appealing in principle, such end-to-end systems pose challenges for high-stakes biomedical extraction, including limited controllability, difficulty enforcing domain-specific constraints, and challenges in attaching verifiable sentence-level evidence. Consequently, OCR-driven pipelines augmented with explicit schemas and provenance tracking remain attractive for applications prioritizing transparency, reproducibility, and auditability [31–33].

Schema-constrained and evidence-centric extraction

A persistent challenge in scientific PDF IE is controlling error propagation at scale. OCR noise can distort drug names and numeric values, layout errors can detach captions or disrupt reading order, and semantic models may hallucinate when evidence is ambiguous. Analyses such as *Document Parsing Unveiled* demonstrate that many downstream failures originate from early-stage structural misalignment rather than limitations in language modeling [34]. This motivates explicit control mechanisms beyond increased model capacity.

Schema-constrained extraction constrains outputs to typed fields with closed vocabularies, explicit negative rules, and well-defined decision hierarchies, reducing spurious inference by limiting the model's degrees of freedom [35, 36]. Complementarily, evidence-centric design requires sentence- or paragraph-level justifications for key decisions, binding each structured value to its textual provenance and enabling systematic audit. These principles align with recent work such as CogDoc, which emphasizes structured, cognitively grounded reasoning over complex documents to improve reliability and interpretability [37].

In summary, despite advances from heuristic parsing to layout-aware transformers and end-to-end neural transcription, the primary challenge for biomedical evidence synthesis is not incremental accuracy, but robust *end-to-end reliability* under heterogeneous layouts and document quality. Schema-constrained extraction with explicit provenance provides a practical foundation for scalable, traceable, and reproducible scientific annotation.

Methods

Overview of the computational pipeline

We developed an end-to-end computational pipeline to transform full-text biomedical articles in PDF format

into structured, analysis-ready records. The system couples a schema-constrained representation of domain knowledge with a large-scale, asynchronous orchestration layer built on top of the Mistral OCR engine. At a high level, the pipeline (i) ingests and normalizes PDF documents, (ii) segments them into page-level chunks to respect model context limits, (iii) submits each chunk to a schema-aware OCR model configured with multiple, domain-specific payloads, (iv) merges and reconciles chunk-level annotations into a single study-level record, and (v) exports tabular and human-readable artifacts for downstream analysis and quality assurance.

The overall design is guided by three principles: (i) *conservatism* (fields are populated only when explicitly supported by the article), (ii) *traceability* (each complex decision is backed by verbatim text from the source manuscript), and (iii) *scalability* (hundreds of long PDFs can be processed in parallel while respecting upstream rate limits).

An overview of the full pipeline architecture is shown in Fig. 3, and the complete implementation is available on GitHub [38].

Algorithm 1 formalizes the high-level end-to-end workflow, while Algorithm 2 details the bounded parallel OCR execution and retry logic used within each document.

Require: Corpus \mathcal{D} ; payload schemas \mathcal{P} ; max pages k ; concurrency C ; rate ρ ; retry budget R_{\max} ; backoff (b_{\min}, b_{\max})

Ensure: Study-level records \mathcal{R} and structured exports

```

1: Load persistent index  $\mathcal{I}$ 
2: Discover PDFs  $\mathcal{F}$  in  $\mathcal{D}$ 
3: Initialize semaphore sem with capacity  $C$ 
4: Initialize rate limiter rl with rate  $\rho$ 
5:  $\mathcal{R} \leftarrow []$ 
6: for all  $f \in \mathcal{F}$  do
7:    $h \leftarrow \text{SHA1}(\text{CanonicalID}(f))$ 
8:   if  $\text{IsProcessed}(\mathcal{I}, h)$  then
9:     continue
10:  end if
11:   $u \leftarrow \text{ToBase64DataURL}(f)$ 
12:   $n \leftarrow \text{GetPageCount}(f)$ 
13:  if  $\text{InvalidPageCount}(n)$  then
14:     $\text{LogAndSkip}(f)$ 
15:    continue
16:  end if
17:   $\mathcal{U} \leftarrow \text{BuildDocumentUnits}(n, k, u)$ 
18:   $\mathcal{A} \leftarrow \text{ProcessDocumentUnits}(u, \mathcal{U}, \mathcal{P}, \text{sem}, \text{rl}, R_{\max}, b_{\min}, b_{\max})$ 
19:   $r \leftarrow \text{IntegrateResults}(\mathcal{A}, h)$ 
20:   $\text{ExportArtifacts}(u, r, h)$ 
21:   $\text{UpdateIndex}(\mathcal{I}, h)$ 
22:  Append  $r$  to  $\mathcal{R}$ 
23: end for
24:  $\text{ExportCorpusTables}(\mathcal{R})$ 

```

Algorithm 1 Schema-constrained OCR pipeline for scalable evidence extraction

```

1: function PROCESSDOCUMENTUNITS(  $u, \mathcal{U}, \mathcal{P}, \text{sem}, \text{rl}, R_{\max}, b_{\min}, b_{\max}$ )
2:   Initialize mapping  $\mathcal{A}$ 
3:   for all  $x \in \mathcal{U}$  in parallel do
4:     for all  $p \in \mathcal{P}$  in parallel do
5:        $\mathcal{A}[p][x] \leftarrow \text{BoundedOCRCall}(u, x, p, \text{sem}, \text{rl}, R_{\max}, b_{\min}, b_{\max})$ 
6:     end for
7:   end for
8:   return  $\mathcal{A}$ 
9: end function

10: function BOUNDEDOCRCALL(  $u, x, p, \text{sem}, \text{rl}, R_{\max}, b_{\min}, b_{\max}$ )
11:    $r \leftarrow 0, b \leftarrow b_{\min}$ 
12:   while  $r \leq R_{\max}$  do
13:     Acquire(sem)
14:     WaitRateLimit(rl)
15:      $(ok, a, err) \leftarrow \text{MistralOCR}(u, x, p)$ 
16:     Release(sem)
17:     if  $ok$  then
18:       return  $a$ 
19:     end if
20:     if IsRetryable(err) then
21:       Sleep( $b$ )
22:        $b \leftarrow \min(2b, b_{\max})$ 
23:        $r \leftarrow r + 1$ 
24:     else
25:       return MarkFailed( $x, p, err$ )
26:     end if
27:   end while
28:   return MarkFailed( $x, p, \text{retry\_exhausted}$ )
29: end function

```

Algorithm 2 Bounded parallel OCR execution and retry logic

Document corpus and ingestion

The corpus consists of biomedical research articles in PDF format, including randomized trials, observational studies, pharmacokinetic studies, diagnostic test accuracy studies, and systematic reviews, all pertaining to direct oral anticoagulant (DOAC) therapy and DOAC level measurement. The pipeline assumes that PDFs are available in a local or network-mounted directory, but it makes no assumptions about journal, publisher, or layout.

As illustrated in Fig. 4, the ingestion module performs deterministic file discovery, hashing, and resume-aware initialization (i.e., the pipeline consults a persistent index of previously processed documents and skips any whose source key already appears in the output, enabling incremental execution without redundant OCR calls) before any downstream processing is invoked.

Each document undergoes a standard ingestion procedure:

1. Canonical identification. For each PDF, we compute a stable source key as $h = \text{SHA1}(\text{CanonicalID}(f))$,

where CanonicalID denotes a deterministic normalization of the file identifier (Unicode normalization, case-folding, path separator normalization, and whitespace collapsing) before hashing. This hash serves as a stable surrogate identifier (*source key*) used to index all downstream artefacts (tables, charts, and markdown reconstructions). The source key provides reproducible linking between extracted records and their originating PDFs while avoiding exposure of file paths or directory structures. Stability is guaranteed within a processing environment and across repeated runs, but not across arbitrary renaming unless the same canonical identifier is preserved.

2. Normalization to base64. The full binary content of each PDF is read and encoded as a base64 string. This string is subsequently wrapped in a standard `data:application/pdf;base64,...` URL and passed to the OCR engine as the primary document payload. This representation avoids

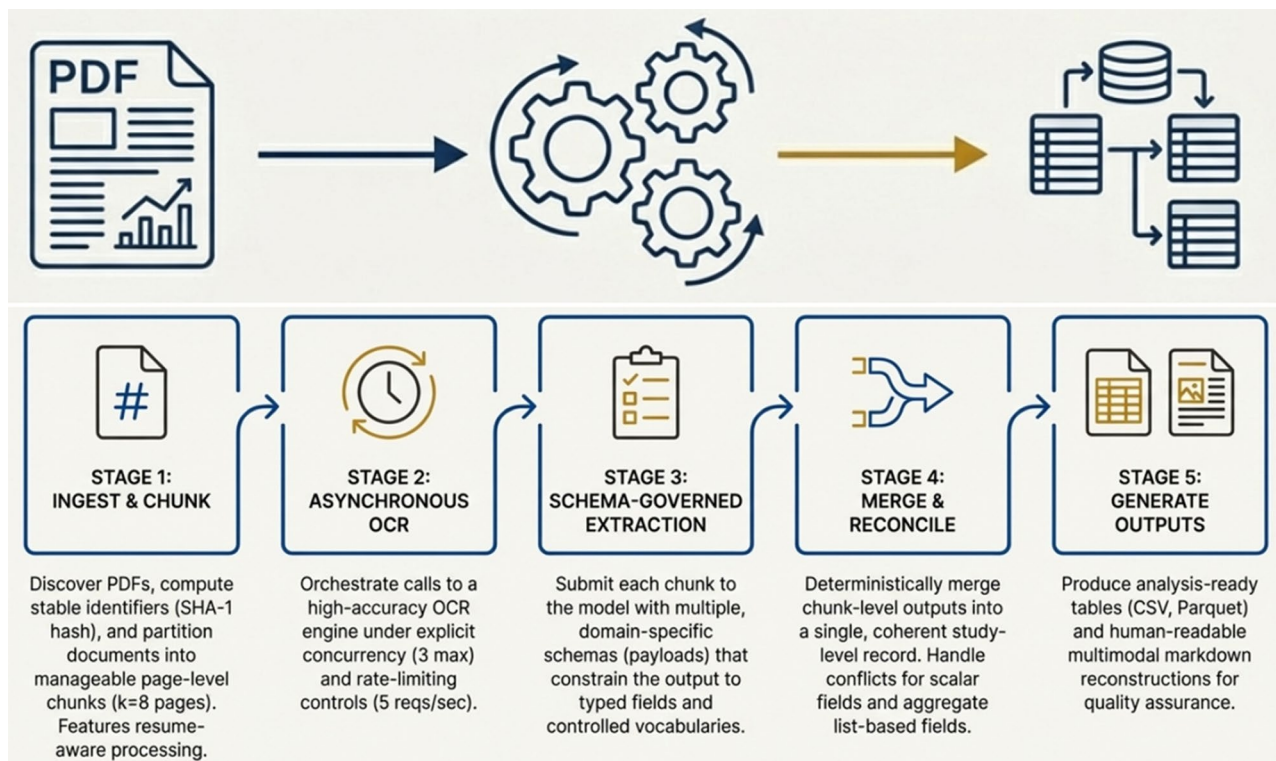


Fig. 3 Overview of the schema-constrained OCR pipeline. PDFs are ingested and chunked, processed asynchronously with schema-constrained OCR, merged into study-level records, and exported as structured tables and reviewable outputs

intermediate file uploads and ensures consistent handling across platforms.

3. Page-level introspection. The number of pages in each document is computed using a PDF parser. Documents with zero or invalid page counts

are excluded from further processing. The page count (denoted n) is later used to construct non-overlapping page chunks.

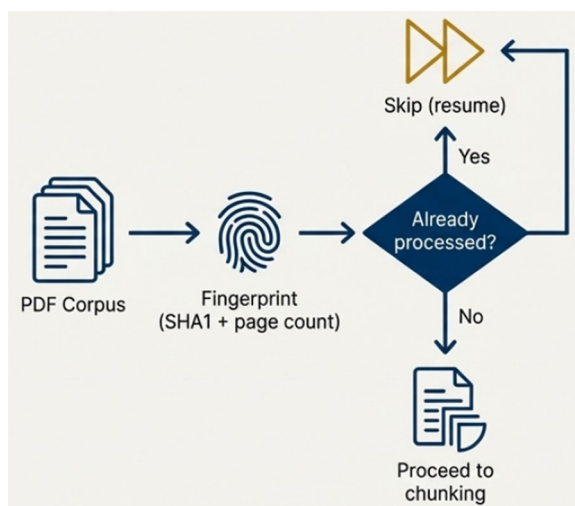


Fig. 4 Overview of the ingestion and resume logic. The system discovers PDF files, computes file identifiers and page counts, encodes documents, and uses a persistent index to determine whether each file requires processing. Documents already present in the index are skipped to ensure efficient, incremental execution

The ingestion layer maintains a persistent index of already processed documents. When the pipeline is re-run on an updated or expanded corpus, this index is consulted, and documents whose source key already appears in the final annotation table are skipped. This resumable execution pattern is crucial for large-scale evidence synthesis workflows where corpora evolve.

Schema-constrained AI inference design

To constrain the behaviour of the OCR model and align extraction with domain expertise, we constructed a set of structured schemas representing the variables required for downstream evidence synthesis. These schemas were implemented using typed data models (Pydantic) but are best understood as formal ontologies comprising controlled vocabularies, decision hierarchies, and explicit negative rules.

Five major payloads were defined, each capturing a coherent domain:

1. bibliographic and study design characteristics;
2. patient population, indications, and subgroups;
3. laboratory methods and assay-related descriptors;

Table 1 Overview of schema payload blocks and the high-level constructs they encode

Payload	Conceptual domain	Representative constructs
Meta-design	Bibliography and study design	Journal and article metadata; field/specialty; year; first-author affiliation; primary study design (e.g., randomized controlled trial, cohort, diagnostic test accuracy); verbatim design justification.
Population & indications	Patient population and clinical context	Total patients with level measurements; DOAC molecules assayed (Apixaban, Rivaroxaban, Edoxaban, Betrixaban, Dabigatran); indications for anticoagulation (e.g., atrial fibrillation, VTE); explicitly studied subgroups (e.g., CKD, bariatric surgery, obesity, urgent surgery); indications for DOAC level measurement (e.g., guide thrombolysis, confirm adherence).
Methods	Assays and pre-analytical variables	DOAC level measurement methods (e.g., LC-MS/MS, calibrated anti-Xa assays, ecarin-based assays, qualitative point-of-care tests); assay descriptors and synonym mapping; pre-analytical conditions (tube type, centrifugation, storage); concurrent conventional and global coagulation testing (PT, aPTT, thrombin generation, viscoelastic tests).
Outcomes	Timing, thresholds, and clinical outcomes	Timing of sample collection relative to dosing (peak, trough, random, not reported); thresholds for DOAC concentration and their use in clinical management; turnaround time; whether clinical outcomes were measured; outcome types (bleeding, thromboembolism); follow-up duration bands; formal outcome definitions (e.g., ISTH, BARC).
Diagnostic performance	Performance of surrogate assays	Diagnostic metrics (sensitivity, specificity, PPV, NPV) for categorical cutoffs; continuous correlations (Spearman, Pearson) between surrogate assays and DOAC levels; comparator assay families (PT, aPTT, TT, dTT, LMWH-calibrated anti-Xa, viscoelastic tests, thrombin generation).

4. timing, thresholds, clinical outcomes, and follow-up; and
5. diagnostic performance metrics for surrogate assays.

Each payload combines: (i) primitive fields (e.g., integer patient counts, verbatim text), (ii) categorical fields restricted to a closed set of labels using enumerations, and (iii) parallel “evidence” fields that store the exact sentence(s) or paragraph(s) from which each classification decision was derived.

Payload blocks and target constructs

Table 1 summarizes the five payload blocks and highlights the conceptual constructs they target.

All payloads share a common guideline philosophy: the model is explicitly instructed not to infer or guess, to set fields to `null` when information is ambiguous or absent, to rely exclusively on explicit statements in the main body of the PDF (excluding references and acknowledgments), and to resolve conflicts only when a clearly dominant statement can be identified (e.g., a formal definition in the [Methods](#) section).

For complex categorical decisions, the schema descriptions encode explicit hierarchies. For example, diagnostic test accuracy studies must be recognized as such and not mislabelled as generic prospective cohorts, and pharmacokinetic studies must not be conflated with broader clinical outcome cohorts when the primary objective is exposure quantification. These rules are expressed directly in the field-level documentation consumed by the OCR model, thereby aligning the model’s decision logic with the expectations of human experts.

Sentence-level evidence for auditability

For each high-level classification (e.g., study design, relevant subgroup, indication for DOAC level measurement, threshold used for clinical management, outcome

definition), there is an associated “sentence from text” field. During extraction, the model is requested to quote the exact sentence(s) or paragraph(s) that justify the chosen label(s). This parallel evidence field is critical for post-hoc auditing and explainability, allowing human reviewers to verify whether the model’s interpretation is faithful to the source.

Schema construction and prompt design

The five extraction schemas were developed iteratively through collaboration between NLP engineers and a domain expert in DOAC level measurement. Initial schema drafts were informed by the variables required for downstream evidence synthesis and refined over multiple review cycles against pilot extraction outputs from representative articles spanning diverse study designs and reporting styles.

Each schema payload is accompanied by a natural-language prompt (system message) that instructs the model on extraction behavior. These prompts encode: (i) conservatism rules directing the model to populate fields only when explicitly supported by text and to default to `null` under ambiguity; (ii) conflict resolution hierarchies specifying that formal definitions in [Methods](#) sections take precedence over informal mentions in Discussion or Introduction; and (iii) evidence-quoting requirements mandating verbatim sentence extraction for all high-level classification decisions.

The schema definitions (implemented as Pydantic models) serve dual roles: they constrain the model’s output format through JSON schema enforcement, and they document the extraction logic through detailed field-level descriptions that function as embedded annotation guidelines. This design ensures that schema evolution is self-documenting and that any change to extraction behavior is traceable to a specific schema or prompt revision.

Mistral OCR engine and schema-constrained annotation

The extraction pipeline employs Mistral OCR 3 (API identifier: `mistral-ocr-2512`, released December 2025) as its OCR backbone, a multimodal vision-language model that processes document pages as images and produces structured outputs. Mistral OCR 3 achieves 88.9% accuracy on handwriting recognition and 96.6% on table extraction benchmarks, and generates Markdown output enriched with HTML-based table reconstruction. The model was selected for four reasons: (i) native schema-constrained output mode allowing direct JSON schema enforcement without post-hoc parsing, (ii) high-quality page-level OCR with layout awareness for multi-column scientific PDFs, (iii) a stable commercial API with rate-limiting support suitable for large-scale batch processing, and (iv) competitive pricing (\$2 per 1,000 pages) enabling corpus-scale deployment. Because Mistral OCR 3 is accessed as a commercial API, results may vary across API versions; we report the specific model identifier to support reproducibility.

All downstream interpretation is governed by explicit schema constraints. Rather than producing unstructured text alone, the OCR stage is coupled with schema-aware

annotation that yields fully typed, structured JSON objects conforming to predefined extraction payloads.

The OCR and schema-guided extraction workflow (Fig. 5) provides a high-level schematic of the pipeline in which PDF documents are partitioned into page-level chunks, scheduled asynchronously under explicit concurrency and rate-limiting constraints, and processed through schema-constrained OCR to yield structured, schema-conformant annotations for downstream consolidation.

Document representation and page chunking

Each PDF is presented to the OCR engine as a base64-encoded `data:application/pdf` URL together with a list of page indices to be processed. To respect context limits and control request sizes, documents are partitioned into non-overlapping chunks of at most k pages (with $k = 8$ in the reference configuration). For an n -page article, the set of chunks is:

$$C = \{ \{0, \dots, k-1\}, \{k, \dots, 2k-1\}, \dots \},$$

with the last chunk potentially shorter. Each chunk is processed independently by the OCR model, but all resulting annotations are later merged back into a single study-level record.

Schema-aware OCR requests

For each page chunk, multiple payloads are submitted concurrently to the Mistral OCR engine. Conceptually, the request consists of:

- a *document input* (the page-range-restricted PDF encoded as a base64 URL);
- a *document annotation format*, which exposes to the model the desired schema (e.g., the meta-design payload or the outcomes payload), including the field names, types, and descriptions; and
- an optional *bounding-box annotation format* used when image-level annotations are requested, describing the type and textual explanation for each detected image region (graph, table, text-only image, other).

The OCR model returns an object comprising: (i) page-level text reconstructions (including layout-aware markdown for each page) and (ii) a *document annotation* field which encodes the schema-constrained extraction. Because the schema includes closed vocabularies and detailed descriptions, the response adheres closely to domain expectations. For instance, DOAC molecules are reported as lists constrained to a fixed set of drug names, and outcome definitions are restricted to recognized taxonomies such as the International Society on Thrombosis

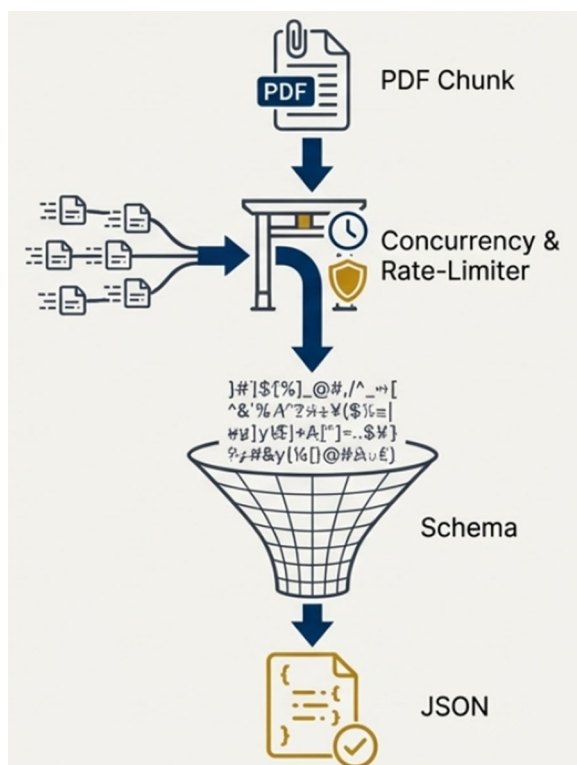


Fig. 5 Schema-constrained OCR extraction. Page-level OCR calls operate under rate-limiting with retry logic, producing structured outputs aligned with predefined extraction schemas. The system identifies captions, extracts text and figures, and compiles chunk-level annotations for downstream merging

and Haemostasis (ISTH) and the Bleeding Academic Research Consortium (BARC).

The same page chunk is independently annotated under each payload, producing complementary but consistent views of the underlying text (e.g., one payload focuses on methods and assays, another on indications and subgroups, another on outcomes and performance metrics).

Image-level annotations and caption-aware chunking

When image analysis is enabled, the pipeline augments textual extraction with a dedicated image-annotation stream. For each page, the Mistral OCR engine identifies all embedded images and returns associated image objects consisting of: (i) a base64 representation of the cropped image region, (ii) a structured bounding-box coordinate set, and (iii) a natural-language description summarizing the visual content (for example “Kaplan–Meier curve for major bleeding”, “Study flow diagram”, or “Box–plot of anti–Xa levels by body weight category”).

Critically, the pipeline also treats *image captions* as independent semantic units. Each caption is isolated from the page layout, extracted as a standalone text chunk, and processed by the OCR model with the same schema-governed constraints applied to body paragraphs. This design acknowledges that captions in biomedical literature often contain substantive methodological or numerical information (e.g., definition of thresholds, subgroup labels, assay descriptions, or outcomes referenced in figures). By elevating captions to first-class textual objects, the pipeline ensures that caption-encoded concepts can contribute to structured extraction when they align with schema definitions, while still preserving the caption’s provenance for later auditing.

Image-level outputs contribute to the pipeline in two ways: (i) *caption-aware text extraction*, where captions are processed as first-class text units that contribute directly to structured extraction when they contain schema-relevant information (e.g., threshold definitions, assay descriptions, or outcome labels), and (ii) *multi-modal markdown reconstructions*, where image metadata and base64-encoded figures are embedded alongside text to support efficient human review. However, the pipeline does not perform figure or plot data extraction (e.g.,

reading numeric values from bar charts or Kaplan–Meier curves); this is noted as future work. The combination of caption-aware text extraction and image metadata lays the groundwork for future multimodal extensions such as figure classification, quantitative data extraction from plots, or transformer-based fusion models that integrate textual and visual evidence.

Asynchronous orchestration and large-scale processing

To achieve high throughput while respecting external constraints, the pipeline orchestrates OCR calls using an asynchronous execution model with explicit rate-limiting and retry policies.

Concurrency and rate-limiting

Two distinct mechanisms control the pace of interaction with the OCR service:

1. Global concurrency control. A fixed upper bound on the number of simultaneously active OCR requests is enforced using a semaphore. This ensures that neither the remote service nor the local system is overwhelmed by too many in-flight requests.
2. Time-based rate-limiting. An additional rate-limiter enforces a minimum inter-request interval, effectively capping the number of OCR calls per second. Before each OCR call is dispatched, the rate-limiter checks the time elapsed since the previous request and, if necessary, introduces a short delay to maintain the desired requests-per-second (RPS) level.

Robust retry strategy

Transient API errors (e.g., HTTP 429 responses, “rate limit” or “quota” messages) are handled via an exponential backoff strategy. When such an error is detected, the pipeline:

- classifies the exception as rate-limit-related based on the status code and error text;
- suspends subsequent calls for a brief period; and
- retries the failed request up to a small, predefined maximum number of attempts (three in the reference configuration), doubling the wait time

Table 2 Key system-level parameters controlling document chunking, concurrency, and throughput

Parameter	Default value	Rationale
Maximum pages per request (<i>k</i>)	8	Balances context utilization and robustness; small enough to keep each OCR call well within model context limits and to localize failures to a small part of the document.
Maximum concurrent OCR calls	3	Limits simultaneous requests to the OCR service, preventing saturation while still providing meaningful parallelism on typical multi-core systems.
OCR request rate (RPS)	5	Caps the number of calls per second to respect upstream rate-limit guidance and allow steady-state throughput without frequent throttling.
Maximum retry attempts	3	Provides a safety margin for transient failures while avoiding long-running retry loops in the presence of persistent errors.

after each failure within specified minimum and maximum bounds.

This ensures that short-lived spikes in demand or temporary service throttling do not cause large-scale failures, while also preventing unbounded retry loops that might exacerbate upstream load.

System-level configuration parameters

Table 2 summarizes the principal system-level parameters controlling chunking, concurrency, and throughput, together with the rationale for their default values.

Through the combination of chunking, bounded concurrency, and rate-limiting, the pipeline can process large corpora stably and predictably, with graceful degradation under adverse network or service conditions.

The principal design choices underlying these parameters were made to balance extraction quality against operational robustness. The chunk size $k = 8$ was empirically chosen to balance context utilization against failure localization: larger chunks risk exceeding the effective context window and entangle failures across many pages, whereas smaller chunks fragment cross-page information and increase API call volume. Non-overlapping, deterministic chunking ensures each page is processed exactly once, avoiding deduplication complexity; cross-chunk information is recovered during the merging phase ([Study-level record construction and tabularization](#) section). The decomposition into five complementary schema payloads reflects the domain structure of DOAC-related evidence synthesis and keeps individual payloads within manageable complexity; running all five payloads

on each chunk ensures comprehensive variable coverage without requiring the model to attend to all constructs simultaneously.

Study-level record construction and tabularization

Each page chunk yields a set of structured annotations for each payload. To obtain a single, coherent record per study, these annotations must be reconciled across chunks and payloads.

As illustrated in Fig. 6, schema-conformant annotations extracted from individual page chunks are reconciled into a single study-level record using deterministic merging rules that enforce consistency for scalar fields, aggregate list-valued fields, and preserve sentence-level evidence across chunks.

Within-payload merging across chunks

For a given payload (e.g., methods or outcomes), the annotations from all page chunks belonging to the same PDF are merged as follows:

- For fields representing unique study-level properties (e.g., publication year, primary study design), the pipeline requires that all non-`null` values be identical. If conflicting values are detected, the record is flagged for manual review rather than arbitrarily selecting one.
- For multi-valued fields (e.g. lists of DOAC molecules, subgroups, thresholds, comparator assays), values across chunks are combined into a set and then converted back to a list to remove duplicates while preserving order of first appearance.

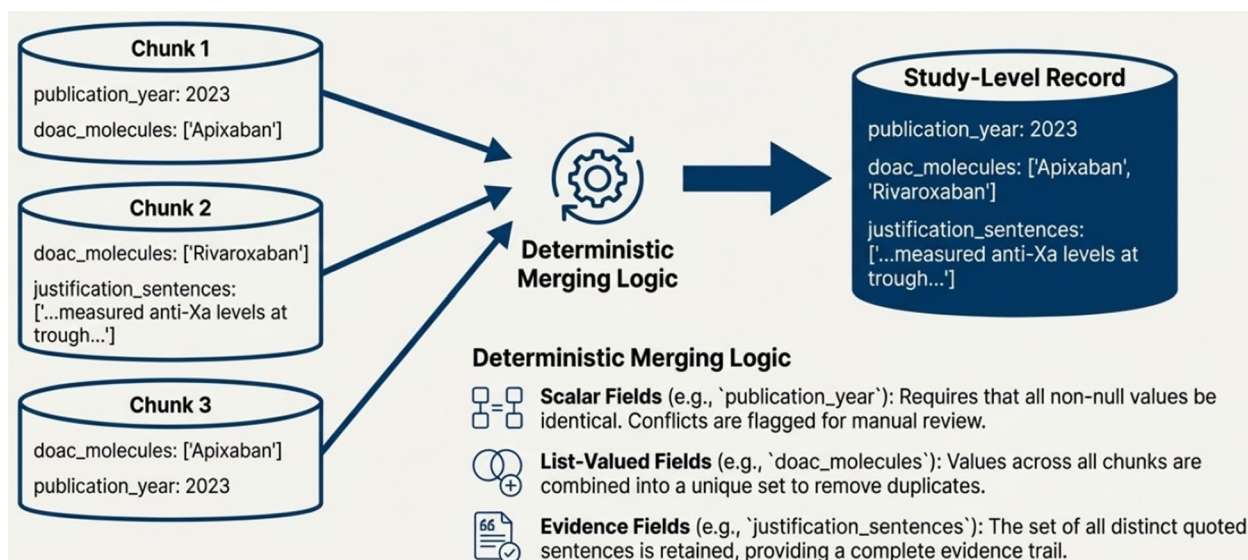


Fig. 6 Deterministic consolidation of chunk-level annotations into a study-level record. Outputs extracted independently from page-level chunks are merged using schema-defined rules: scalar fields must agree across chunks, list-valued fields are de-duplicated and combined, and sentence-level evidence is retained for auditability

- For evidence fields (sentence-level justifications), the set of all distinct quoted sentences is retained, enabling reviewers to see every place in the article where a particular concept was described.

In this way, the system produces a payload-specific, study-level annotation that integrates information scattered across the entire article.

Cross-payload integration

Once each payload has been merged across chunks, the resulting payload-level annotations are further integrated into a single study-level record. Because the schemas are designed to be complementary rather than overlapping, conflicts across payloads are rare. Where multiple payloads touch on related constructs (e.g., laboratory methods appearing in both methods and diagnostic performance payloads), the integration logic preserves all non-redundant values.

The final study-level records are exported in both row-oriented (CSV) and columnar (Parquet) formats. The ordering and naming of columns are derived directly from the schema definitions, yielding a stable, machine-generated data dictionary and ensuring that every row is schema-conforming even when individual studies do not populate all fields.

Post-extraction synthesis, quality assurance, and visualization

Beyond producing structured tables, the pipeline generates artefacts to support exploratory analysis and quality assurance.

Markdown reconstructions

For each document, a human-readable markdown reconstruction is produced by combining:

1. a header summarizing the structured annotations for that document (payload-specific key-value pairs rendered as bolded text); and
2. page-level markdown derived from the OCR output, in which image placeholders are replaced with base64-encoded inline images and their corresponding annotations.

These reconstructions provide a compact, navigable view of each article and allow reviewers to quickly verify whether extracted values (e.g., study design, thresholds, subgroups) are consistent with the text and figures.

Aggregation and distributional summaries

At the corpus level, multi-valued fields (e.g., lists of DOACs, subgroups, tests) are normalized through Unicode case-folding, whitespace collapsing, and mapping of

known synonyms to canonical labels (e.g., “anti-Xa assay”, “chromogenic anti-FXa”, and “calibrated anti-Xa” all map to “Calibrated anti-Xa assay”). After normalization, fields are, when appropriate, combined into composite stratified variables (such as DOAC molecule by type of coagulation test). Value-count distributions are computed on the canonical labels for all fields and exported as a multi-sheet spreadsheet in which each sheet corresponds to one variable. For each variable, the sheet contains both a complete frequency table and an automatically generated bar chart summarizing the most prevalent categories.

This artifact provides an immediate, visual overview of the extracted dataset and facilitates the detection of anomalous patterns (e.g, unexpectedly frequent labels, missing categories in the presence of textual evidence) that may warrant targeted evaluation of the schemas or model behaviour.

Proxy transcription quality assessment

Because character-level ground truth is typically unavailable at corpus scale, OCR fidelity was assessed using a proxy transcription-quality score rather than direct character error rate. A rule-based validation module computed document- and chunk-level quality indicators from OCR outputs, including: (i) *character corruption*, measured as the fraction of tokens containing three or more consecutive non-alphanumeric characters; (ii) *numeric sanity*, checking whether extracted patient counts, thresholds, and percentages fall within plausible clinical ranges (e.g., patient counts > 0 , percentages ≤ 100); (iii) *schema conformance*, the fraction of typed fields that pass Pydantic validation without coercion errors; and (iv) *token density*, the ratio of extracted tokens to expected tokens per page based on page dimensions and detected layout type. These indicators were aggregated into a bounded proxy score used for operational monitoring and failure detection rather than formal OCR evaluation.

Throughput is reported in terms of *OCR API requests*, where a single request corresponds to one page chunk or caption unit processed under a specified schema payload. End-to-end processing time was measured per document as the wall-clock interval from dispatch of the first page-level OCR request to completion of all payload annotations and study-level consolidation.

Implementation considerations, reproducibility, and extensibility

The pipeline is controlled via a small set of configuration variables that govern model credentials, concurrency, image-annotation behaviour, and overwrite policies. These parameters allow the same codebase to be deployed across environments ranging from personal workstations to high-throughput compute clusters while maintaining consistent behaviour.

Several design choices explicitly support reproducibility and extensibility:

- Deterministic schema-derived columns. All output columns are generated from the schema definitions, ensuring that the data model is versioned and that downstream analyses can be reproduced exactly when the same schema and corpus are used.
- Stable source identifiers. Hash-based source keys provide a robust linkage between structured records and their originating PDFs without relying on mutable file paths or filenames.
- Sentence-level justifications. Parallel evidence fields mean that every high-level classification can be traced back to specific sentences in the source article, enabling independent adjudication and calibration against human annotators.
- Incremental processing. The resume mechanism, combined with append-only tabular storage, allows existing corpora to be expanded without repeating computationally expensive OCR operations, and without compromising the integrity of previously annotated records.
- Domain-agnostic architecture. The pipeline separates *domain-agnostic architectural components*—chunking, asynchronous orchestration, merging logic, provenance tracking, and export—from *domain-specific components*—schema definitions, controlled vocabularies, and prompt instructions. Although the present schemas are tailored to DOAC level measurement, adapting the system to a new biomedical domain (e.g., oncology biomarkers, prognostic scores, imaging modalities) requires replacing only the schema definitions and vocabularies; the pipeline infrastructure remains unchanged. In our experience, developing a new domain schema requires iterative collaboration between an NLP engineer and a domain expert over approximately 2–4 weeks, depending on reporting heterogeneity and the complexity of the target variables.

In combination, these features yield a robust, generalizable framework for large-scale, schema-guided extraction of structured evidence from full-text biomedical literature using modern OCR–LLM technology.

Results

Corpus coverage and end-to-end execution

The pipeline processed the full corpus in a single end-to-end run without manual intervention. In total, $N_{pdf} = 734$ PDFs comprising 7,228 pages were discovered, ingested, and assigned stable source identifiers via SHA-1 hashing of canonicalized file identifiers. The

resume-aware index prevented redundant processing: $N_{skip} = 2$ PDFs already present in the final annotation table were detected and skipped, and all remaining documents proceeded through chunking, OCR, structured extraction, merging, and export.

To characterize corpus heterogeneity and document length, the pipeline recorded per-document page counts at ingestion and used these to drive deterministic chunk formation. Using a maximum of $k = 8$ pages per OCR request, the corpus yielded $N_{chunks} = 978$ non-overlapping page chunks. Chunk sizes were stable by construction (median = 8 pages; shorter final chunks occurred only when the page count was not divisible by k). In addition to page chunks, caption isolation produced $N_{cap} = 824$ caption-level text units (derived from detected figure and table captions) that were processed under the same schema constraints. Captions were treated as first-class semantic objects to improve recall for constructs frequently expressed outside narrative text (e.g., assay descriptors, threshold definitions, subgroup labels, and outcome references).

Throughput, service constraints, and call volume

The pipeline sustained stable throughput under the reference configuration (concurrency cap of three, rate limit of five RPS). The observed steady-state throughput was $R = 3.1$ OCR requests per second; the gap between target and observed RPS reflects backpressure from bounded retries, caption-level processing, and heterogeneous chunk processing times.

Each page chunk was annotated with five schema payloads, yielding $N_{chunk-req} = 978 \times 5 = 4,890$ payload-specific requests. Caption units contributed $N_{cap-req} = 824 \times 5 = 4,120$ additional requests. The full corpus therefore required $N_{req} = 9,010$ schema-constrained OCR requests, excluding retries.

Under the reference configuration, the mean processing time was $T_{mean} = 11$ seconds per PDE, increasing monotonically with document length. Resume-aware indexing converts the pipeline into an incremental workflow, such that re-processing an expanded corpus invokes OCR only for newly discovered or modified documents.

Proxy transcription quality and operational OCR robustness

On a stratified sample of documents spanning layout complexity (multi-column text, dense tables) and input quality (born-digital versus scanned), the median proxy transcription score (Proxy transcription quality assessment section) was $Q_{median} = 98.6\%$ with a central 95% interval of [97.8%, 99.2%]. This estimate should be interpreted as a deployment-oriented quality signal sensitive to error modes likely to affect downstream schema

extraction, not as a replacement for character-level OCR evaluation.

Operational robustness was further supported by automatic recovery from transient failures. Across all requests, the pipeline recorded $N_{retry} = 10$ retry events; all were resolved within the bounded exponential back-off policy, and none propagated to fatal document-level failures.

Evaluation protocol for extraction correctness

Extraction correctness was assessed through expert review on a simple random sample of 50 studies drawn from the full corpus. The evaluation focused on synthesis-critical constructs for which extraction errors would materially affect downstream evidence synthesis, including study design, assay classification, timing of measurement relative to dosing, outcome definitions, follow-up duration, and use of clinical thresholds.

For each sampled study, two reviewers collaboratively (not blinded, see [Discussion](#) section) examined the structured record alongside its sentence-level evidence fields and the corresponding multimodal markdown reconstruction. A field was scored as correct if (i) the extracted value matched the source document and (ii) the associated evidence explicitly supported the value without reliance on inference. For categorical variables, correctness required agreement with schema-defined labels; for numeric variables, correctness required concordant values and units as reported in the manuscript, allowing for trivial formatting differences.

Disagreements and ambiguous cases were resolved by joint re-examination of the cited evidence spans within the reconstruction. This protocol is intended to assess practical end-to-end extraction fidelity under schema and provenance constraints rather than benchmark performance against a fully annotated gold-standard corpus. Accordingly, reported correctness reflects the reliability of extracted variables in realistic evidence synthesis settings rather than competitive model performance.

Ablation-style evaluation of schema constraints and refinement

This analysis evaluates schema constraints as an AI inference control mechanism, assessing their effect on stability, auditability, and resistance to over-inference under long-document processing. Schema-constrained extraction demonstrated high operational stability and controllability across multi-page, multi-chunk documents. During within-payload consolidation, scalar fields were required to agree across all non-null chunk-level outputs. Any disagreement triggered an explicit conflict flag rather than automated resolution. Across the full corpus, such conflicts were infrequent, indicating that typed schemas, closed vocabularies, and explicit decision

constraints substantially reduced instability introduced by page-level chunking. For list-valued fields, set-based aggregation was applied to merge redundant mentions across chunks while preserving the order of first appearance, producing consistent study-level representations even when constructs were reported repeatedly across sections, tables, or captions.

The inclusion of sentence-level evidence fields materially improved auditability and error containment. For schema components configured to require textual justification, explicit supporting sentences were populated in >99% of cases where the construct was present and extractable from OCR output. Instances of missing evidence were predominantly attributable to information reported exclusively within dense tables or complex figures, where relevant text could not be reliably isolated as sentence-level OCR output. In these cases, the system defaulted to conservative behavior, leaving fields unpopulated or flagging records for review rather than inferring unsupported values. This conservative, evidence-gated behavior is an intentional design choice: in biomedical evidence synthesis, omission with explicit traceability is preferable to producing seemingly complete but weakly grounded outputs, as clinical and methodological rigor prioritize fidelity to source evidence over forced completeness.

To assess the impact of iterative refinement, we conducted expert review on a random sample of 50 studies drawn from the corpus. Evaluation focused on synthesis-critical fields for which extraction errors would materially affect downstream analysis, including study design, assay classification, timing of measurement relative to dosing, outcome definitions, follow-up duration, and threshold usage. Fields were scored as correct only when (i) the extracted value matched the source document and (ii) the associated evidence field explicitly supported the decision without reliance on inference. Fields not reported in the source document were treated as ineligible rather than incorrect.

Table 3 summarizes expert-assessed correctness before and after iterative schema and prompt refinement. Improvements were most pronounced for constructs that are typically dispersed across narrative text, tables, and captions, including outcome definitions, follow-up duration, clinical outcome presence, and clinical threshold usage. These gains reflect the cumulative effect of clarifying schema definitions, tightening negative rules that prevent over-inference, and expanding controlled vocabularies to better capture domain-specific variation. Correctness estimates are reported with Wilson 95% confidence intervals to reflect uncertainty arising from the finite evaluation sample and to provide appropriate coverage for proportions near 0 or 1. For several domains, no errors were observed in the reviewed sample following

Table 3 Effect of schema and prompt refinement on expert-assessed extraction correctness. Values are reported as percentage correct with Wilson 95% confidence intervals, based on expert review of a random sample of $n = 50$ studies per component

Schema component	Before	After
Clinical outcome definitions and follow-up duration	33% (20.8–45.8)	95% (86.5–98.9)
Presence/type of clinical outcomes (bleeding, VTE, etc.)	50% (36.6–63.4)	100% (92.9–100.0)
Relevant patient subgroups (CKD, bleeding, reversal, bariatric surgery)	63% (50.1–75.9)	95% (86.5–98.9)
Indications for DOAC level measurement	70% (56.2–80.9)	90% (78.6–95.7)
Assay types used for DOAC level measurement (LC–MS, dTT, calibrated anti–Xa)	70% (56.2–80.9)	95% (86.5–98.9)
Conventional coagulation tests reported (PT, aPTT, TT, etc.)	70% (56.2–80.9)	90% (78.6–95.7)
Global coagulation testing (TGA, TEG/ROTEM)	70% (56.2–80.9)	100% (92.9–100.0)
Timing of DOAC level measurement relative to dosing (peak, trough, random)	70% (56.2–80.9)	90% (78.6–95.7)
Study design classification	75% (62.6–85.7)	100% (92.9–100.0)
Use of thresholds for clinical management	75% (62.6–85.7)	100% (92.9–100.0)
Pre-analytical variables (tube type, centrifugation, storage, freezing)	80% (67.0–88.8)	95% (86.5–98.9)
Indications for anticoagulation (AF, VTE, etc.)	90% (78.6–95.7)	95% (86.5–98.9)
Reported numeric DOAC concentration thresholds	100% (92.9–100.0)	100% (92.9–100.0)

Wilson confidence intervals were computed using a two-sided normal critical value of $z = 1.96$. Percentages correspond to the proportion of eligible fields scored as correct. Rows with 100% correctness indicate no observed errors in the reviewed sample and should not be interpreted as definitive upper bounds on performance

refinement; given the evaluation size, these results should be interpreted as an absence of observed errors rather than as definitive upper bounds on performance.

Overall, these results indicate that schema-constrained extraction, when combined with explicit provenance requirements and iterative expert-in-the-loop refinement, yields stable and controllable behavior under full-text, multi-chunk processing. Importantly, performance gains were achieved not through increased model capacity, but through tighter specification of target constructs and constraints, reinforcing the role of explicit schemas and evidence linkage as primary drivers of practical extraction reliability in biomedical evidence synthesis.

Record completeness and reporting heterogeneity

We report completeness to characterize both extraction coverage and underlying reporting heterogeneity. Completeness was defined as the proportion of studies in which a component contained at least one non-null value *supported by evidence when evidence fields were available*. Because component-level completeness can be inflated when a single easy field is present, we report both payload-level completeness (Table 4) and field-level

Table 4 Payload-level completeness rates across the structured schema. Completeness is the proportion of studies with at least one non-null value in the payload, with evidence used to confirm support where applicable

Schema component	Completeness (%)
Bibliographic metadata	99.9%
Study design classification	99.7%
Population and indications	91.6%
Methods and assay details	98.6%
Outcomes and follow-up descriptions	97.9%
Diagnostic performance metrics	98.2%
Evidence fields (where eligible)	99.9%

completeness for representative high-difficulty constructs that are central to synthesis.

At the payload level, bibliographic metadata and study design fields were populated for nearly all documents, reflecting both the salience of these constructs and their frequent explicit reporting in abstracts and headers. Population-level completeness was lower, driven by heterogeneity in reporting of patient counts with level measurements, subgroups, and measurement indications. Methods and assay fields showed high completeness because assay families and test names are typically stated explicitly in [Methods](#) sections and captions. Outcomes and diagnostic performance payloads exhibited high payload-level completeness but also substantial within-payload sparsity, consistent with the fact that many studies report only subsets of clinical outcomes, follow-up windows, or performance metrics.

To mitigate misinterpretation of payload-level completeness, we additionally inspected field-level completeness for difficult constructs that are often omitted or inconsistently phrased (e.g., formal outcome taxonomy, explicit follow-up duration, timing relative to dosing, and threshold usage). These constructs showed substantially lower completeness than the payload aggregates, consistent with heterogeneous reporting practices and underscoring the value of evidence-linked records for targeted review.

The global pattern of extraction coverage and sparsity is visualized in the missingness matrix (Fig. 7), which highlights dense coverage for bibliographic/design fields and structured sparsity in specialized outcome and performance constructs.

Corpus-level distributions and derived stratifications

The structured records support automated corpus-level characterization without additional manual




Fig. 7 Missingness matrix for extracted fields. Rows are studies (sorted by completeness), columns are schema fields (grouped by payload: Meta–design, Population/Indications, Methods, Outcomes, Diagnostic Performance). Light cells indicate missing (null) values; dark cells are populated. Coverage is highest for bibliographic/design fields, while outcome and performance fields show more missingness, reflecting reporting heterogeneity. Right marginal bar shows per-study completeness

Distribution of DOACs Included (> 3)

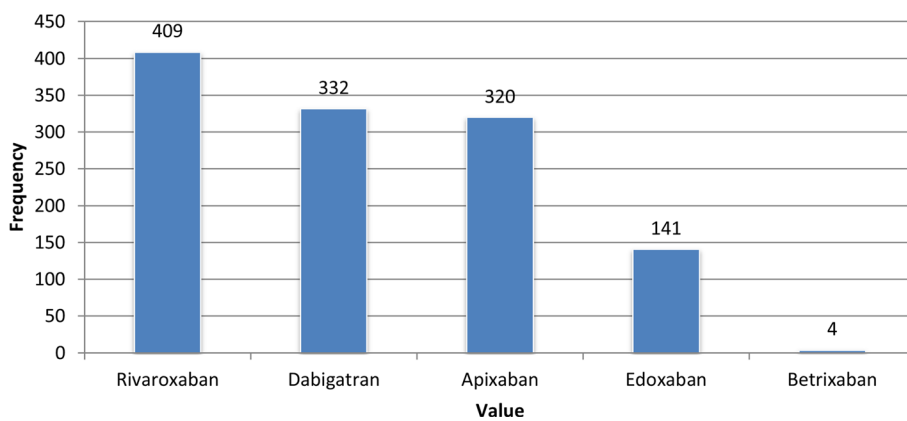


Fig. 8 Representative system-generated frequency distribution for DOAC molecules identified by schema-constrained extraction. Frequencies correspond to the number of studies reporting each molecule; analogous summaries are generated automatically for all extracted variables. Only values occurring more than three times are shown

normalization. All list-valued fields were reconstructed into native list structures and used to compute frequency distributions for each categorical variable. The resulting distributions were internally coherent and reflected the temporal availability and historical uptake of individual DOACs rather than contemporary prescribing

prevalence. In particular, dabigatran, which was the first DOAC introduced into clinical practice, appeared frequently in earlier studies focused on drug level quantification, whereas rivaroxaban and apixaban predominated in later publications. In contrast, edoxaban appeared less frequently, consistent with its later regulatory approval,

and betrixaban was rare, reflecting its limited clinical adoption and subsequent withdrawal from the market. Figure 8 shows a representative system-generated distribution for DOAC molecules; analogous summaries were generated automatically for all schema fields.

To expose methodological heterogeneity, we constructed composite stratifications via Cartesian products between selected population attributes and assay/reporting variables (e.g., DOAC molecule \times concurrent coagulation testing). These stratifications enabled multilevel tabulations that are useful for synthesis planning and QA. For example, among studies measuring apixaban, a majority concurrently reported calibrated anti-Xa assays, while thrombin generation assay parameters appeared only in a small subset of documents. Such derived variables provide immediate, reviewable signals of reporting patterns and potential subgroup structure within the corpus.

Multimodal reconstructions and provenance-linked review artifacts

For each study, the pipeline generates a complete, human-readable markdown reconstruction that integrates OCR-derived page text, caption-level text units, structured annotation summaries, and inline figures encoded as base64 with associated semantic descriptions. This unified artifact enables efficient expert review by allowing extracted variables to be inspected in direct proximity to their narrative, tabular, and caption-level context, without reliance on external document viewers.

Reproducibility was assessed under fixed execution conditions. Given identical document inputs, fixed schema versions, and a fixed OCR model version, all downstream stages—including schema enforcement, chunk-level merging, study-level consolidation, and export to CSV/Parquet and markdown formats—are deterministic. Consequently, repeated executions yield identical serialized outputs and markdown reconstructions, conditional on identical OCR responses. Stable source keys and chunk-index metadata ensure deterministic traceability from each structured field to its originating document and page range, supporting transparent auditing and post-hoc verification.

Error handling, recovery, and end-to-end reliability

The orchestration layer provided bounded recovery from transient failures and ensured that localized errors did not cascade into corpus-level failures. Across the full run, the system encountered $N_{err} = 15$ transient OCR errors (including rate-limit responses and intermittent service errors). All were handled automatically via exponential backoff with a bounded retry budget; no fatal document-level failures occurred. Importantly, the resume-aware index ensured that any interrupted runs could be restarted without repeating completed OCR

work, supporting robust long-running execution on evolving corpora.

Together, these results demonstrate that schema-constrained OCR can be operated as a stable, scalable, and auditable extraction workflow: call volumes and throughput are measurable under service limits, extraction correctness improves with systematic refinement, completeness patterns reflect both extraction capability and true reporting heterogeneity, and the generated artifacts support efficient provenance-linked expert review.

Discussion

This work presents an end-to-end, OCR-driven pipeline for transforming full-text biomedical PDFs into structured, analysis-ready records with explicit schema constraints and sentence-level provenance. The central contribution is not a new document model, but an engineering and methodological pattern that targets the practical failure modes of evidence synthesis: heterogeneous layouts, dispersed reporting across tables and captions, long-document context limits, and the need for auditable outputs that can be reviewed and corrected. Across a corpus focused on DOAC level measurement, the system achieved stable large-scale processing, produced internally consistent study-level records, and yielded measurable gains in extraction fidelity after iterative refinement of prompts and schemas.

Principal findings and practical value

Three findings are most salient for evidence synthesis workflows. First, the pipeline demonstrated reliable corpus-scale execution under strict service constraints through deterministic page chunking, bounded concurrency, and resumable processing. In practice, this shifts automation from a fragile, one-off batch process to an incremental annotation workflow that can be re-run as corpora evolve without re-incurring unnecessary OCR costs. Second, schema-constrained extraction improved controllability: restricting categorical fields to closed vocabularies, enforcing typed outputs, and requiring evidence for high-level decisions reduced spurious inference and enabled systematic post-hoc auditing. Third, the combination of structured outputs with human-readable markdown reconstructions created a dual representation of each study, enabling both quantitative synthesis (CSV/Parquet) and rapid qualitative review (source-linked text, captions, and images in a single artifact). Together, these design choices align automated extraction with the operational needs of systematic reviews, where traceability and error localization are as important as raw accuracy.

Positioning relative to prior scholarly PDF processing

Prior work in scholarly PDF extraction spans rule-based parsing, feature-driven sequence labeling, layout-aware

transformers, and end-to-end neural transcription. Classic systems such as GROBID and CERMINE established robust baselines for recovering document structure and bibliographic fields [5, 6], while modern layout-aware models demonstrated that spatial grounding can substantially improve extraction in visually rich documents [18–20]. More recent efforts revisit transcription itself, converting document images into structured representations and narrowing the boundary between OCR and downstream understanding [26, 27]. Concurrently, TrialMind has demonstrated the use of large language models for accelerating clinical evidence synthesis tasks such as screening and data extraction [39], and recent evaluations of AI-powered screening tools highlight both the promise and the current limitations of automated literature triage [40]. Our work complements these strands by emphasizing end-to-end *reliability under biomedical constraints*. Rather than optimizing for a single benchmark task, we treat the pipeline as a synthesis-oriented system in which upstream parsing, OCR, extraction, merging, and auditing mechanisms must jointly minimize downstream risk. In particular, we highlight that the most consequential errors in evidence synthesis frequently arise from cascading failures across stages (reading order, caption detachment, OCR artifacts, over-inference), a pattern also emphasized in recent analyses of document parsing pipelines [34].

We acknowledge that the current work does not include head-to-head comparisons against alternative extraction systems. Direct comparison was not performed because existing tools target different task formulations (e.g., TrialMind focuses on screening and PICO extraction [39]; GPT-4V and Claude lack native schema-constrained output modes), making fair comparison non-trivial without substantial adaptation. Comparative evaluation against both constrained and unconstrained extraction paradigms is planned as a priority for future work (Discussion section, Future directions).

Benchmark-oriented evaluations were intentionally deprioritized because the primary objective of this work was domain-faithful, end-to-end extraction for biomedical evidence synthesis in a clinically specialized setting rather than optimization on generic document benchmarks. In the context of DOAC level measurement, many synthesis-critical variables—such as assay family, timing relative to dosing, pre-analytical conditions, and outcome definitions—are sparsely reported, expressed heterogeneously, and often embedded in captions or tables that are poorly represented in standard benchmarks. Evaluating the pipeline under realistic full-corpus conditions with domain-specific schemas, controlled vocabularies, and sentence-level provenance, therefore, provides a more meaningful assessment of practical extraction fidelity and failure modes for this use case than isolated

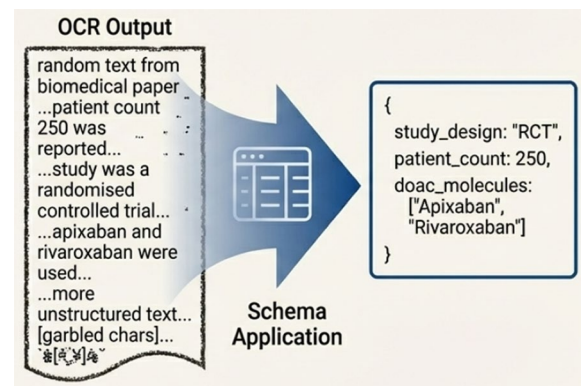


Fig. 9 Schema-constrained mapping of unstructured OCR text to structured, auditable study variables with explicit semantics and provenance

benchmark scores, which we view as complementary but insufficient for high-stakes, domain-constrained evidence synthesis.

In contrast to unconstrained OCR or text-first pipelines, the proposed approach explicitly limits the hypothesis space through typed schemas and evidence gating, reducing silent over-inference under noisy transcription. This distinction is particularly consequential in biomedical evidence synthesis, where plausibly correct but unsupported values impose downstream audit costs.

OCR errors in drug names, numeric values, and units can directly corrupt extracted fields and propagate to downstream synthesis. Schema constraints partially mitigate this risk: closed vocabularies reject unrecognizable drug names, and typed numeric fields reject non-numeric strings. However, subtle OCR errors such as digit substitution (e.g., “150” misread as “l50”) can pass schema validation; the proxy quality score and sentence-level evidence fields provide secondary detection mechanisms by enabling reviewers to trace suspicious values back to source text. Low-quality scans were rare in this predominantly born-digital corpus, but the pipeline’s conservative null-defaulting behavior provides a safeguard: when OCR output is ambiguous, fields are left unpopulated rather than populated with uncertain values.

Why schema constraints and provenance matter in biomedical extraction

Biomedical evidence synthesis imposes requirements that differ from many general document IE settings. Variables such as assay family, dosing-time alignment (peak versus trough), and clinical outcome definitions are often expressed implicitly and dispersed across methods, text, tables, and captions. In this setting, unconstrained generation can produce outputs that appear plausible but are not faithfully grounded in the source document.

Our schema-first approach mitigates this risk by explicitly narrowing the hypothesis space through typed fields and controlled vocabularies, and by requiring

sentence-level evidence for high-level decisions. Figure 9 illustrates this process, showing how noisy OCR output is transformed into normalized, schema-constrained study variables with explicit semantics. The practical value of this design lies in its support for an audit loop, where extracted records can be systematically reviewed, corrected, and used to iteratively refine schemas and prompts. Consistent improvements observed following such refinements reflect increased controllability and clarity of target constructs, rather than benchmark-driven optimization.

Expert-in-the-loop prompt refinement and audit workflow

An expert-in-the-loop refinement process with a DOAC domain specialist was central to improving extraction fidelity for synthesis-critical variables. A random sample of 50 full-text articles was reviewed, with field-level correctness annotated and recurrent failure modes documented, particularly for pre-analytical variables, coagulation testing, patient subgroups, and clinical outcomes. Expert feedback informed iterative schema and prompt revisions, including tighter keyword mappings, explicit “None/NA” gating when outcomes were not measured, and rules to distinguish study focus from background mentions.

This workflow treated errors as specification gaps rather than isolated model failures, enabling systematic convergence toward a stable extraction configuration. Disagreements were localized to individual fields with supporting sentence-level evidence, reviewed by the expert, and translated into updated constraints and prompts that improved robustness under heterogeneous biomedical reporting.

Despite iterative refinement, certain error categories persist. The most common residual errors arise from: (i) information reported exclusively in complex tables whose multi-row and multi-column structure cannot be fully linearized by OCR, (ii) implicit information requiring cross-section reasoning beyond the boundaries of a single page chunk (e.g., an outcome definition in Methods referenced only by shorthand in Results), and (iii) ambiguous terminology where schema-defined controlled vocabularies do not fully capture the reporting variation encountered across journals and clinical specialties. These errors are systematic rather than random, suggesting that further schema refinement and improved table handling represent the most productive mitigation strategies for future iterations.

Implications for scalable and practical evidence synthesis

The pipeline enables a workflow in which structured extraction is not an endpoint, but a substrate for downstream evidence synthesis and quality assurance. The structured outputs are directly actionable within synthesis workflows: sentence-linked records allow reviewers to

rapidly verify study design, assay families, timing relative to dosing, and outcome definitions without re-reading entire manuscripts. By binding extracted variables to explicit textual evidence, the system supports efficient adjudication and reduces the risk of silent errors.

From a practical perspective, the observed end-to-end processing time has direct implications for scalability. Prior methodological studies report that manual extraction of synthesis-critical variables from full-text biomedical articles by trained reviewers requires tens of minutes per study, with mean extraction times of approximately 30–40 minutes in systematic review settings [41]. Under the reference configuration, the proposed pipeline processed a full-text PDF in a mean of 11 seconds. While automated outputs still require expert verification, this represents an order-of-magnitude reduction in primary extraction time, shifting human effort from exhaustive manual abstraction to targeted audit and adjudication.

At the corpus level, automated distributional summaries function as both descriptive analytics and quality assurance signals, revealing schema drift or systematic omissions. Conflict detection for scalar fields prevents silent propagation of inconsistent values, while evidence-gated null values distinguish true reporting absence from extraction failure. Together with machine-readable datasets and human-readable multimodal reconstructions, these features shift the role of human experts from exhaustive manual abstraction toward higher-value synthesis, validation, and interpretation.

Limitations

Several limitations warrant consideration. First, extraction correctness was assessed through internal consistency checks and expert review on limited random samples rather than through statistically powered evaluation against a fully annotated gold-standard corpus or large-scale benchmarks such as SciNLP [42]. Consequently, reported performance gains should be interpreted as evidence of practical, end-to-end extraction fidelity under evidence synthesis constraints rather than as definitive benchmark-level generalization across tasks or competing systems.

Second, expert review was conducted collaboratively without blinding, and no intermediary formal inter-annotator agreement metric (e.g., Cohen's κ) was computed. While this approach was appropriate for iterative refinement, it limits the interpretability of reported correctness rates and should be addressed in future evaluations through blinded, independent annotation with formal agreement assessment.

Third, OCR quality was estimated using downstream validation heuristics and schema-based sanity checks rather than character-level ground truth annotations [43]. While these operational signals are appropriate for

large-scale deployment and failure detection, they may fail to capture subtle transcription errors that do not propagate to schema violations and therefore do not replace formal OCR evaluation.

Fourth, although caption-aware chunking improves recall for information embedded in figures and tables, the pipeline does not yet perform fine-grained alignment between figures, captions, tables, and in-text references, nor does it extract quantitative values directly from plots or complex tables beyond OCR-accessible text [44]. These limitations restrict the full exploitation of multimodal evidence in visually dense scientific articles.

Fifth, the evaluated schemas were developed for DOAC level measurement. While the architectural framework is domain-agnostic, effective transfer to other biomedical domains depends on the availability and quality of domain-specific schema definitions, controlled vocabularies, and expert input. Inadequately specified schemas may limit both extraction accuracy and interpretability in new clinical contexts [45, 46].

Finally, as with any OCR-LLM-based system, performance is influenced by upstream model behavior and service-level constraints. Although the pipeline is architecturally decoupled from any specific OCR provider, changes in model versions, API behavior, or rate limits can affect throughput, latency, and error profiles. These factors underscore the importance of careful orchestration, robust error handling, and cost-aware execution, consistent with prior analyses of cascading failures in document parsing pipelines [34] and emerging work on adaptive validation and cost control in large-scale ML systems [47, 48].

Future directions

Several extensions could further strengthen both scientific rigor and practical utility. First, a larger-scale evaluation against a curated gold-standard subset with blinded expert annotation would enable more formal assessment of extraction performance, including inter-annotator agreement for synthesis-critical constructs such as outcome definitions, assay classification, and timing of measurement relative to dosing. Such an evaluation would help disentangle model limitations from intrinsic ambiguity and heterogeneity in biomedical reporting.

Second, future work should include systematic head-to-head comparisons against alternative extraction paradigms, including text-first pipelines and unconstrained generative approaches. Direct comparison would allow more precise quantification of the empirical benefits of schema constraints and sentence-level provenance, particularly with respect to over-inference, silent errors, and the downstream audit burden imposed on human reviewers.

Methodologically, confidence-aware extraction represents a promising extension. Attaching calibrated

uncertainty estimates at the field level would allow downstream synthesis workflows to weight extracted variables, triage records for manual review, or defer aggregation when confidence falls below predefined thresholds. This would better align automated extraction with evidence synthesis practices that explicitly account for uncertainty and variable reporting quality.

On the document understanding side, richer structural grounding could be achieved by incorporating explicit layout graphs or region-level linking between in-text references, captions, tables, and figures. Improved handling of complex tables and dense numerical reporting would be particularly valuable for laboratory and diagnostic studies, where key quantities are often expressed outside narrative text and are prone to transcription or alignment errors.

Finally, the schema-constrained and provenance-centric design pattern introduced here could be extended beyond extraction to support adjacent stages of evidence synthesis, including semi-automated screening, structured PICO formulation, and risk-of-bias assistance. Extending the framework in this direction would require maintaining auditability, evidence linkage, and conservative inference as first-class constraints, ensuring that any downstream assistance remains transparent, reviewable, and compatible with established evidence synthesis standards. More broadly, such efforts align with emerging calls for a synthesis-ready research ecosystem in which study design, reporting, and publication practices are optimized to facilitate timely and comprehensive evidence integration [49].

Overall, these results indicate that AI-driven, OCR-based structured extraction can be scaled without sacrificing controllability when deterministic chunking, schema constraints, and evidence capture are combined within an auditable workflow. In biomedical evidence synthesis, where the cost of silent extraction errors is high, this work demonstrates an end-to-end AI pipeline that produces reviewable structured records and reproducible multimodal artifacts, enabling systematic inspection, iterative refinement, and reliable large-scale annotation of heterogeneous scientific PDFs.

Conclusion

This work demonstrates that full-text biomedical PDFs can be transformed into structured, analysis-ready evidence at a corpus scale without sacrificing auditability. We present a systems and methodological contribution in the form of an end-to-end, OCR-driven pipeline that integrates (i) deterministic document ingestion with resume-aware hashing, (ii) page-level chunking and asynchronous orchestration under explicit concurrency and rate constraints, (iii) schema-constrained, typed extraction across multiple domain-specific payloads, and (iv) study-level consolidation via conflict-aware merging for

scalar fields and set-based aggregation for list-valued fields. Crucially, the pipeline binds synthesis-critical variables to sentence-level provenance, enabling transparent verification of extracted values against explicit supporting statements in the source manuscripts.

Applied to a large corpus of studies on direct oral anticoagulant (DOAC) level measurement, the system processed all documents without manual intervention while maintaining stable throughput under rate limiting. It produced fully reproducible tabular datasets (CSV and Parquet) alongside multimodal, caption-aware markdown reconstructions that support efficient expert review. Iterative refinement of schemas and prompts yielded substantial gains in practical extraction fidelity for challenging constructs such as outcome definitions, follow-up duration, assay family classification, and identification of global coagulation testing. These findings underscore a central lesson for evidence synthesis automation: reliable end-to-end performance arises not solely from model capacity, but from explicit specification, conservative inference, and enforceable traceability throughout the extraction pipeline.

Although evaluated within a DOAC-focused domain, the architectural pattern is general. Adapting the pipeline to other biomedical areas requires only the replacement of schema definitions and controlled vocabularies, without changes to the underlying orchestration, merging, or auditing logic. By constraining outputs to typed schemas and systematically binding extracted values to sentence-level evidence, the proposed approach provides a practical foundation for scalable and auditable literature annotation. In turn, this enables more reproducible evidence synthesis workflows, faster characterization of reporting heterogeneity across large corpora, and a reduced burden of human review by directing expert attention toward provenance-linked decisions, flagged conflicts, and synthesis-critical uncertainties.

Abbreviations

AF	Atrial fibrillation
aPTT	Activated partial thromboplastin time
BARC	Bleeding Academic Research Consortium
CKD	Chronic kidney disease
CRF	Conditional random field
DOAC	Direct oral anticoagulant
dTT	Diluted thrombin time
IE	Information extraction
IEP	Information extraction pipeline
ISTH	International Society on Thrombosis and Haemostasis
LC-MS/MS	Liquid chromatography-tandem mass spectrometry
LMWH	Low-molecular-weight heparin
OCR	Optical character recognition
PDF	Portable document format
PPV	Positive predictive value
PT	Prothrombin time
RPS	Requests per second
TEG	Thromboelastography
TGA	Thrombin generation assay
TT	Thrombin time
VTE	Venous thromboembolism

Authors' contributions

Pouria Mortezaagha and Arya Rahgozar conceived and designed the model, developed the computational framework, and implemented the end-to-end pipeline. Joseph Shaw provided domain expertise in DOAC level measurement, developed the schema specifications, and led prompt refinements. Bowen Sun contributed to pipeline development, testing, and integration of key processing components. All authors performed the evaluation and analysis. All authors contributed to the interpretation of the results, critically revised the manuscript, and approved the final version.

Funding

This work was supported by the Ottawa Hospital Research Institute.

Data availability

The full-text articles analyzed were obtained from open-access sources. Derived datasets, intermediate outputs, and analysis scripts are available from the corresponding author upon reasonable request, subject to publisher licensing restrictions.

Declarations

Ethics approval and consent to participate

The models used in this study were not trained, fine-tuned, or adapted on the analyzed articles. All document processing was performed in an inference-only setting, and no extracted content was used for model training or retained beyond the analytical scope of this study. The corpus consisted exclusively of open-access full-text articles obtained from publicly available sources, and no proprietary, restricted, or patient-identifiable data were used. The pipeline operates solely on published scientific literature and does not introduce new personal data, human subjects, or privacy risks beyond those inherent to the original publications.

Consent for publication

Not applicable.

Competing interests

The authors declare no competing interests.

Author details

¹Methodological Implementation Research, Ottawa Hospital Research Institute, Ottawa, ON, Canada

²Inflammation and Chronic Disease, Ottawa Hospital Research Institute, Ottawa, ON, Canada

³School of Engineering Design and Teaching Innovation, University of Ottawa, Ottawa, ON, Canada

Received: 31 December 2025 / Accepted: 26 March 2026

Published online: 14 April 2026

References

1. Tóth B, Berek L, Gulácsi L, Péntek M, Zrubka Z. Automation of systematic reviews of biomedical literature: a scoping review of studies indexed in PubMed. *Syst Rev*. 2024;13(1):174.
2. Büchter RB, Weise A, Pieper D. Development, testing and use of data extraction forms in systematic reviews: a review of methodological guidance. *BMC Med Res Methodol*. 2020;20(1):259. <https://doi.org/10.1186/s12874-020-01143-3>.
3. Affengruber L, van der Maten MM, Spiero I, Nussbaumer-Streit B, et al. An exploration of available methods and tools to improve the efficiency of systematic review production: a scoping review. *BMC Med Res Methodol*. 2024;24(1):210. <https://doi.org/10.1186/s12874-024-02320-4>.
4. Mikriukov A, Senokosov A, Succi G, Tormasov A, Plaksin Y, Trofimova E, et al. AI tools for automating systematic literature reviews. In: Proceedings of the 2025 International Conference on Software Engineering and Computer Applications. New York: ACM; 2025. pp. 25–30.
5. Lopez P. GROBID: Combining automatic bibliographic data recognition and term extraction for scholarship publications. In: Agosti M, Borbinha J, Kapidakis S, Papatheodorou C, Tsakonias G. (editors) Research and Advanced Technology for Digital Libraries. ECDL 2009. Lecture Notes in Computer

- Science, vol 5714. Springer, Berlin: Heidelberg; 2009. https://doi.org/10.1007/978-3-642-04346-8_62.
6. Tkaczyk D, Szostek P, Fedoryszak M, Dendek PJ, Bolikowski L. Cermine: automatic extraction of structured metadata from scientific literature. *International Journal on Document Analysis and Recognition (IJ DAR)*. 2015;18:317–35. <https://doi.org/10.1007/s10032-015-0249-8>.
 7. Isaac Councill, C. Lee Giles, Min-Yen Kan. ParsCit: an Open-source CRF Reference String Parsing Package. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco. European Language Resources Association (ELRA). 2008. <https://aclanthology.org/L08-1291/>. Accessed 15 Dec 2025.
 8. Shafait F, Keyzers D, Breuel TM. Document Image Analysis Using the Documentum Algorithm. In: *International Workshop on Document Analysis Systems (DAS)*. 2008.
 9. Kim Y, Gotoh K, Toyosada M. Automatic two-dimensional layout using a rule-based heuristic algorithm. *J Mar Sci Technol*. 2003;8(1):37–46.
 10. Prasad A, Kaur M, Kan MY. Neural parsit: a deep learning-based reference string parser. *Int J Digit Libr*. 2018;19(4):323–37.
 11. Lo K, Wang LL, Neumann M, Kinney R, Weld DS. S2ORC: The Semantic Scholar Open Research Corpus. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*. 2020. pp. 4969–4983. <https://doi.org/10.18653/v1/2020.acl-main.447>.
 12. Clark C, Divvala S. PDFFigures 2.0: Mining Figures from Research Papers. In: *Proceedings of the 16th ACM/IEEE-CS Joint Conference on Digital Libraries (JCDL)*. 2016. <https://doi.org/10.1145/2910896.2910904>.
 13. Safder I, Hassan SU, Visvizi A, Noraset T, Nawaz R, Tuab S. Deep learning-based extraction of algorithmic metadata in full-text scholarly documents. *Inf Process Manag*. 2020;57(6):102269.
 14. Shen Z, Lo K, Wang LL, Kuehl B, Weld DS, Downey D. VILA: improving structured content extraction from scientific PDFs using visual layout groups. *Trans Assoc Comput Linguist*. 2022;10:376–92. <https://doi.org/10.1162/tacl.2022.10.0376>.
 15. Zhong X, Tang J, Yepes AJ. PubLayNet: Largest Dataset Ever for Document Layout Analysis. 2019. arXiv preprint [arXiv:1908.07836](https://arxiv.org/abs/1908.07836).
 16. Li M, Xu Y, Cui L, Huang S, Wei F, Li Z, et al. DocBank: A Benchmark Dataset for Document Layout Analysis. In: *Proceedings of the 28th International Conference on Computational Linguistics (COLING)*. 2020. <https://aclanthology.org/2020.coling-main.82/>. Accessed 15 Dec 2025.
 17. Ouyang L, Qu Y, Zhou H, Zhu J, Zhang R, Lin Q, et al. OmniDocBench: Benchmarking Diverse PDF Document Parsing with Comprehensive Annotations. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2025.
 18. Xu Y, Li M, Cui L, Huang S, Wei F, Zhou M. LayoutLM: Pre-training of Text and Layout for Document Image Understanding. 2019. arXiv preprint [arXiv:1912.13318](https://arxiv.org/abs/1912.13318).
 19. Xu Y, Xu Y, Lv T, Cui L, Wei F, Wang G, et al. LayoutLMv2: Multi-modal Pre-training for Visually-Rich Document Understanding. 2020. arXiv preprint [arXiv:2012.14740](https://arxiv.org/abs/2012.14740).
 20. Appalaraju S, Jasani B, Kota BU, Xie Y, Manmatha R. DocFormer: End-to-End Transformer for Document Understanding. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 2021. https://openaccess.thecvf.com/content/ICCV2021/papers/Appalaraju_DocFormer_End-to-End_Transformer_for_Document_Understanding_ICCV_2021_paper.pdf. Accessed 15 Dec 2025.
 21. Wang D, Raman N, Sibue M, Ma Z, Babkin P, Kaur S, et al. DocLLM: A Layout-Aware Generative Language Model for Multimodal Document Understanding. In: *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL) (Long Papers)*. 2024.
 22. Chen Y, Liu R, Zheng J, Wen D, Peng K, Zhang J, et al. Graph-based Document Structure Analysis. In: *International Conference on Learning Representations (ICLR)*. 2025.
 23. Beltagy I, Lo K, Cohan A. SciBERT: A Pretrained Language Model for Scientific Text. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics; 2019. pp. 3613–3618. <https://aclanthology.org/D19-1371/>. Accessed 15 Dec 2025.
 24. Rostam KZR, Kertész G. Advances in pre-trained language models for domain-specific text classification: a systematic review. *ACM Trans Intell Syst Technol*. 2025;16(6):1–41.
 25. Zhang Q, Ding K, Lv T, Wang X, Yin Q, Zhang Y, et al. Scientific large language models: a survey on biological & chemical domains. *ACM Comput Surv*. 2025;57(6):1–38.
 26. Blecher L, Cucurull G, Scialom T, Stojnic R. Nougat: Neural Optical Understanding for Academic Documents. 2023. arXiv preprint [arXiv:2308.13418](https://arxiv.org/abs/2308.13418).
 27. Liu Y, Yang B, Liu Q, Li Z, Ma Z, Zhang S, et al. TextMonkey: An OCR-Free Large Multimodal Model for Understanding Document. 2024. arXiv preprint [arXiv:2403.04473](https://arxiv.org/abs/2403.04473).
 28. Zhang Y, Liu H, Wang Z. Neural OCR for Complex Scientific Documents. 2024. [arXiv:2409.16934](https://arxiv.org/abs/2409.16934).
 29. Poznanski J, Henderson L, Soldaini L, et al. olmOCR: Unlocking Trillions of Tokens in PDFs with Vision Language Models. 2025. arXiv preprint [arXiv:2502.18443](https://arxiv.org/abs/2502.18443).
 30. Kim G, Hong TH, Yim M, Nam J, Park J, Kim J, et al. OCR-Free Document Understanding Transformer. In: *European Conference on Computer Vision (ECCV)*. 2022. [arXiv:2111.15664](https://arxiv.org/abs/2111.15664).
 31. Gururaja S, Zhang Y, Tang G, Zhang T, Murphy K, Yi YT. Collage: Decomposable Rapid Prototyping for Co-Designed Information Extraction on Scientific PDFs. In: Ghosal T, Mayr P, Singh A, Naik A, Rehm G, Freitag D, editors. *Proceedings of the Fifth Workshop on Scholarly Document Processing (SDP 2025)*. Vienna: Association for Computational Linguistics; 2025. pp. 72–82. <https://doi.org/10.18653/v1/2025.sdp-1.7>.
 32. Schäfer H, Schmidt CS, Wutzkowsky J, Lorek K, Reinartz L, Rückert J, et al. A Multimodal Pipeline for Clinical Data Extraction: Applying Vision-Language Models to Scans of Transfusion Reaction Reports. 2025. [arXiv:2504.20220](https://arxiv.org/abs/2504.20220).
 33. Sinha R, S RB. Digitization of Document and Information Extraction using OCR. 2025. [arXiv:2506.11156](https://arxiv.org/abs/2506.11156).
 34. Zhang Q, Wang B, Huang VSJ, Zhang J, Wang Z, Liang H, et al. Document Parsing Unveiled: Techniques, Challenges, and Prospects for Structured Information Extraction. 2025. [arXiv:2410.21169](https://arxiv.org/abs/2410.21169).
 35. Hassani-Pak K, Singh A, Brandizi M, Hearnshaw J, Parsons JD, Amberkar S, et al. KnetMiner: a comprehensive approach for supporting evidence-based gene discovery and complex trait analysis across species. *Plant Biotechnol J*. 2021;19(8):1670–8.
 36. Bhattacharya C, Som S. AISAC: An Integrated multi-agent System for Transparent, Retrieval-Grounded Scientific Assistance. 2025. [arXiv:2511.14043](https://arxiv.org/abs/2511.14043).
 37. Xu Q, Wang H, Liu C, Lin F, Chen W. CogDoc: Towards Unified Thinking in Documents. 2025. [arXiv:2512.12658](https://arxiv.org/abs/2512.12658).
 38. Mortezaagha P. Schema-constrained OCR pipeline for auditable biomedical evidence extraction. GitHub repository. 2025. <https://github.com/pouriamr/mistral-ocr-pipeline>. Accessed 20 Dec 2025.
 39. Jin Q, Wang Z, Yang Y, et al. Accelerating clinical evidence synthesis with large language models. *Nat Med*. 2025. <https://doi.org/10.1038/s41591-024-03233-x>.
 40. Ruan M, Fan J, Liu M, Meng Z, Zhang X, Zhang C. Artificial intelligence for the science of evidence synthesis: how good are AI-powered tools for automatic literature screening? *BMC Med Res Methodol*. 2025;25(1):199. <https://doi.org/10.1186/s12874-025-02644-9>.
 41. Sercombe J, Bryant Z, Wilson J. Evaluating a customized version of ChatGPT for systematic review data extraction in health research: development and usability study. *JMIR Form Res*. 2025;9:e68666.
 42. Duan D, Zhang Y, Peng J, Zhang C. SciNLP: A Domain-Specific Benchmark for Full-Text Scientific Entity and Relation Extraction in NLP. 2025. [arXiv:2509.07801](https://arxiv.org/abs/2509.07801).
 43. van Strien D, Beelen K, Ardanuy M, Hosseini K, McGillivray B, Colavizza G. Assessing the impact of OCR quality on downstream NLP tasks. In: *Proceedings of the 12th International Conference on Agents and Artificial Intelligence*. SCITEPRESS - Science and Technology Publications; 2020.
 44. Sarto S, Cornia M, Cucchiara R. Image Captioning Evaluation in the Age of Multimodal LLMs: Challenges and Future Perspectives. 2025. [arXiv:2503.14604](https://arxiv.org/abs/2503.14604).
 45. Rath A. Structured Prompting and Feedback-Guided Reasoning with LLMs for Data Interpretation. 2025. [arXiv:2505.01636](https://arxiv.org/abs/2505.01636).
 46. Chhetri TR, Chen Y, Trivedi P, Jarecka D, Haobsh S, Ray P, et al. STRUCTSENSE: A Task-Agnostic Agentic Framework for Structured Information Extraction with Human-In-The-Loop Evaluation and Benchmarking. 2025. [arXiv:2507.03674](https://arxiv.org/abs/2507.03674).
 47. Ghafouri S, Razavi K, Salmani M, Sanaee A, Botran TL, Wang L, et al. [Solution] IPA: Inference Pipeline Adaptation to achieve high accuracy and cost-efficiency. *J Syst Res*. 2024;4(1). <https://doi.org/10.5070/sr34163500>.
 48. Tu D, He Y, Cui W, Ge S, Zhang H, Shi H, et al. Auto-Validate by-History: Auto-Program Data Quality Constraints to Validate Recurring Data Pipelines. 2023. [arXiv:2306.02421](https://arxiv.org/abs/2306.02421).
 49. Bannach-Brown A, Rackoll T, Macleod MR, et al. Building a synthesis-ready research ecosystem: fostering collaboration and open science to accelerate biomedical translation. *BMC Med Res Methodol*. 2025;25(1):66. <https://doi.org/10.1186/s12874-025-02524-2>.
- ## Publisher's Note
- Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.