An End-to-End Pipeline for Handwritten Prescription Understanding: OCR–NLP Integration and Error Propagation Analysis

OPEN ACCESS

Volume: 13

Special Issue: 2

Month: January

Year: 2026

E-ISSN: 2582-0397

P-ISSN: 2321-788X

Citation:
M, Sindhana Devi, and Hari Prakash K. "An End-to-End Pipeline for Handwritten Prescription Understanding: OCR–NLP Integration and Error Propagation Analysis." *Shanlax International Journal of Arts, Science and Humanities*, vol. 13, no. 2, 2026, pp. 123–33.

DOI:
https://doi.org/10.34293/sijash.v13iS2-i4-Jan.10593

Dr. Sindhana Devi M
Data Science
Kumaraguru College of Liberal Arts and Science, Coimbatore, Tamil Nadu, India

Hari Prakash K
Data Science
Kumaraguru College of Liberal Arts and Science, Coimbatore, Tamil Nadu, India

Abstract
Interpreting handwritten medical prescriptions is a constant challenge in healthcare workflows. This often leads to patient confusion, medication errors, and delays in receiving proper treatment. Even with improvements in digital health systems, reliably understanding prescriptions is limited by differences in handwriting quality and the complexity of medical language. This paper introduces a complete pipeline for understanding handwritten prescriptions. It combines Optical Character Recognition (OCR) with clinical Natural Language Processing (NLP) to turn unstructured prescription images into organized medical information. The proposed framework focuses on system robustness by examining how errors from OCR affect downstream NLP-based medication extraction. To ensure a realistic and repeatable evaluation while respecting privacy, the system is tested using publicly available handwritten text and clinical NLP benchmark datasets. The experimental analysis looks at error variability, performance in entity-level extraction, and how the integration of OCR and NLP affects prescription interpretation. The results show that while OCR and clinical NLP perform reasonably on their own, their interaction is vital for overall system reliability. These findings underline the need for joint evaluation and error-conscious design in automated prescription processing systems. This supports the creation of better, patient-focused digital healthcare solutions.
Keywords: Handwritten Prescription Analysis, Optical Character Recognition, Clinical Natural Language Processing, Medical Entity Extraction, Error Propagation Analysis, Digital Healthcare Systems, Prescription Digitization.

Introduction

Handwritten medical prescriptions still play an important role in healthcare, especially in outpatient clinics, small hospitals, and areas with limited resources. Even with the rise of electronic health record systems, many regions still rely on manual prescription writing. This creates problems with legibility, interpretation, and accessibility. Patients often find it hard to read handwritten prescriptions, which can cause medication errors, delays in treatment, and increased reliance on pharmacists or caregivers for help.

Automating the understanding of prescriptions is challenging due to two main sources of variability. First, handwritten prescriptions vary widely in writing styles, abbreviations, and layouts, making it hard to extract text accurately. Second, even when the text is recognized, the content often uses specialized medical language that requires context to identify medication names, dosages, and instructions correctly. Mistakes at any point can affect patient safety and system usability.

Recent developments in Artificial Intelligence have led to improvements in both Optical Character Recognition (OCR) and Natural Language Processing (NLP). OCR has gotten better at recognizing handwritten and scanned documents. Clinical NLP methods have shown they can effectively extract structured medical information from unstructured text. However, most current methods assess these components separately, reporting OCR accuracy or entity extraction performance without looking at how errors affect the overall process. In real healthcare systems, OCR and NLP work one after the other. Errors in text recognition can significantly impact later medical interpretations.

Additionally, evaluating automated prescription systems realistically is limited by the lack of publicly available end-to-end prescription datasets because of privacy, ethical, and legal issues. Many studies use small proprietary datasets or synthetic examples, which restrict reproducibility and make it hard to evaluate real-world effectiveness. There is a need for evaluation frameworks that use publicly available benchmarks while accurately reflecting the variability of prescriptions and the complexity of medical language.

This paper tackles these issues by presenting an end-to-end pipeline for understanding handwritten prescriptions that combines OCR and clinical NLP within a single framework. Instead of proposing new algorithms for recognition or extraction, it focuses on evaluating the system as a whole, looking at how variability in OCR affects the extraction of medical entities and the overall understanding of prescriptions. Public handwritten text datasets and clinical NLP benchmarks are used to ensure reproducibility, and error propagation analysis is employed to gain insights relevant to patient-focused healthcare workflows.

The contributions of this work are threefold. First, it provides a structured pipeline that links handwritten text recognition with clinical entity extraction for prescription understanding. Second, it introduces a method of evaluation that considers errors and analyses the interaction between OCR and NLP components rather than viewing them separately. Third, it offers practical insights into the limitations and reliability of automated prescription processing, helping to improve the design of more dependable, patient-centred digital healthcare systems.

Background

Handwritten medical prescriptions are still essential in many healthcare settings, such as outpatient clinics, small hospitals, and places with limited resources. Even as electronic health record systems become more common, writing prescriptions mostly relies on manual methods. This is especially true in areas where digital resources are lacking or where quick patient care is a priority. Relying on handwritten prescriptions brings ongoing issues with legibility, interpretation, and access to vital medical information.

From the patient's viewpoint, pharmacists or healthcare professionals often need to interpret handwritten prescriptions. This creates dependence and raises the chances of misunderstandings. Mistakes in interpreting prescriptions can lead to incorrect medication use, wrong dosages, or delays in treatment. These risks underscore the need for dependable systems to digitize and interpret prescriptions, connecting handwritten instructions to information that patients can readily understand.

Recent breakthroughs in Artificial Intelligence have greatly improved document understanding tasks, particularly through Optical Character Recognition (OCR) and Natural Language Processing (NLP). OCR has become better at recognizing handwritten text across various writing styles, while clinical NLP has shown it can effectively extract structured medical information from unstructured text. However, applying these improvements to real-world prescription understanding systems is still difficult due to the wide variety of handwriting styles and the intricacies of clinical language.

One significant limitation in current research is the fragmented evaluation of OCR and NLP tools. Many studies aim to improve OCR accuracy or medical entity extraction by themselves, without looking at how mistakes made during text recognition affect later clinical interpretation. In practical prescription processing, OCR and NLP work together in a sequence. Errors made early in the process can lead to bigger mistakes later on. This interaction is crucial in healthcare, where even small errors can have serious consequences.

A further significant issue in prescription analysis research is the shortage of publicly available, complete prescription datasets. Privacy laws, ethical issues, and legal restrictions limit the sharing of actual prescription images with verified medical annotations. As a result, many proposed systems are tested using private datasets or small pilot studies, which restricts reproducibility and makes it hard to compare studies. This limitation highlights the need for evaluation methods that use publicly available benchmarks while capturing realistic variations in handwriting and clinical language.

Related Work

A. Handwritten Text Recognition for Medical Documents

Handwritten Text Recognition (HTR) has been studied extensively in the area of document digitization. Early methods focused on statistical and rule-based approaches, while more recent studies have used deep learning techniques. Convolutional Neural Networks (CNNs) combined with recurrent models like Long Short-Term Memory (LSTM) networks and Connectionist Temporal Classification (CTC) loss have shown better performance on general handwriting benchmarks. However, interpreting handwritten medical prescriptions remains difficult due to cursive writing styles, inconsistent spacing, abbreviations specific to the field, and frequent overlap between text regions.

Several studies have examined OCR techniques for prescription images and reported moderate recognition accuracy, even when using deep learning models. These studies note that handwriting on prescriptions varies greatly compared to standard handwriting datasets, and there is often insufficient annotated training data. Consequently, OCR performance on prescriptions is very sensitive to the quality of handwriting and the layout of the document. Most prior studies evaluate OCR accuracy alone, without considering how recognition mistakes impact subsequent medical interpretation tasks.

B. Clinical Natural Language Processing for Medication Extraction

Clinical Natural Language Processing (NLP) has been widely used to extract structured information from unstructured medical text. This is especially true for tasks like recognizing medication names, extracting dosages, and identifying administration frequencies. Early systems relied on rule-based methods and medical lexicons. Later work introduced machine learning models such as Conditional Random Fields and neural sequence labelling architectures.

Shared tasks and benchmark datasets in clinical NLP have led to notable progress in extracting medication entities. Recent transformer-based models have achieved high precision and recall on annotated clinical notes. Despite these advancements, the performance of extraction varies by entity type. Medication names are generally easier to identify than dosage and frequency information. Additionally, most clinical NLP systems work under the assumption of clean text derived from electronic health records. Their ability to handle noisy or error-prone OCR outputs has not been thoroughly examined.

C. Integrated OCR–NLP Pipelines in Healthcare Applications

In actual healthcare settings, OCR and NLP components are usually part of an integrated processing pipeline. A limited number of studies have looked into end-to-end document understanding systems that merge image pre-processing, OCR, and NLP for medical documents, such as clinical reports, discharge summaries, and prescriptions. These works highlight the need for optimizing the entire pipeline, including image pre-processing and text normalization after OCR, to enhance extraction accuracy.

However, existing integrated approaches often focus on system design or overall accuracy without examining how OCR and NLP stages interact. Specifically, the impact of OCR errors on clinical entity extraction has not been studied enough. This issue is crucial in prescription analysis, where small recognition errors can result in incorrect medication identification or dosage interpretation, potentially putting patient safety at risk.

D. Evaluation Challenges and Dataset Limitations

A continuing challenge in prescription analysis research is the lack of publicly available, end-to-end datasets containing handwritten prescription images matched with verified medical annotations. Privacy regulations and ethical concerns limit the release of real prescription data, leading many studies to depend on proprietary datasets or small-scale experiments. This lack of standardized datasets hampers reproducibility and makes it hard to compare different studies.

To overcome these limitations, recent research has used benchmark-guided evaluation strategies. These strategies leverage publicly available handwriting datasets for OCR analysis and clinical NLP benchmarks for entity extraction. While these datasets do not capture the full prescription workflows, they allow for controlled and reproducible evaluation of individual pipeline components and enable a systematic analysis of robustness and variability.

E. Positioning of the Present Work

Previous research has made notable advances in handwritten text recognition and clinical medication extraction when these issues are examined separately. However, relatively few studies take a system-level view to investigate how OCR and NLP components work together in an end-to-end prescription understanding pipeline. In particular, the effect of OCR variability on the performance of clinical NLP has not been sufficiently studied. This work aims to fill this gap by focusing on the integration of OCR and NLP and analysing how errors propagate for understanding handwritten prescriptions. The emphasis is not on proposing new recognition or extraction models but on realistic evaluation using public benchmarks. This approach seeks to understand how errors at the component level affect the overall reliability of the system. This perspective is essential for designing robust, patient-centred prescription digitization systems that are suitable for real-world healthcare use.

System Architecture and Methodology

A. System Overview

The proposed system is a complete framework for understanding handwritten prescriptions automatically. It combines Optical Character Recognition (OCR) and clinical Natural Language Processing (NLP). The architecture consists of three main components: Frontend, Backend, and Processing Models. This setup allows for modular processing and organized evaluation. The Frontend offers a simple interface for uploading handwritten prescription images and displaying structured prescription outputs. The Backend oversees the processing workflow. It coordinates image pre-processing, OCR execution, NLP-based entity extraction, and result aggregation. The Processing Models layer includes the OCR module for recognizing handwritten text and the clinical NLP module for extracting medication-related entities. This separation aids reproducibility and allows for detailed system-level analysis.

B. Image Acquisition and Pre-processing

Prescription images come from scanned documents or photographs taken with mobile devices. To minimize variability from lighting, noise, and resolution, a pre-processing pipeline is used. This includes grayscale normalization, noise reduction, contrast improvement, and resizing. These steps enhance OCR stability while keeping handwritten features intact.

C. Handwritten Text Recognition

The OCR module converts the pre-processed prescription images into raw text output. A general handwritten text recognition method is used to mirror real-world conditions. OCR performance is tested with publicly available handwriting benchmarks, measuring Character Error Rate (CER) to capture recognition variability. The OCR output goes directly to the NLP stage without any postprocessing corrections to maintain realistic error behaviour.

D. Clinical NLP-Based Entity Extraction

The recognized text is analysed with clinical NLP techniques to extract important prescription elements, including medication names, dosage details, and administration frequency. Entity extraction is treated as a sequence labelling task and measured against standard clinical NLP benchmarks. Performance is evaluated using precision, recall, and confusion analysis at the entity level.

E. Error Propagation Analysis

A key part of the methodology is examining how errors move through the OCR-NLP pipeline. OCR error metrics are linked to the performance of entity extraction to determine how recognition mistakes affect prescription understanding. This system-level analysis identifies types of entities that are especially sensitive to errors from the OCR stage.

F. Output Representation

The final output of the system is a structured representation of prescription information. This design improves readability and supports patient-focused healthcare applications. The structured output also allows for a quantitative assessment of the system's reliability and durability.

Experimental Setup

A. Hardware and Software Environment

All experiments took place on a local development system with a 2.6 GHz 6-core Intel Core i7 processor, 16 GB DDR4 memory (2667 MHz), and Intel UHD Graphics 630, running macOS Sequoia (version 15.6.1). The software stack was built on Python 3 and included a web-based prescription analysis pipeline. It integrated OCR for extracting handwritten text and clinical NLP for finding medication-related entities. This setup represents a realistic deployment environment for understanding prescriptions on regular hardware.

B. Model Configuration

The prescription understanding pipeline has two modules that work in sequence: a handwritten text recognition module and a clinical NLP module. The OCR component processes pre-processed prescription images and provides raw text output without correction, which maintains realistic recognition errors. The clinical NLP component works directly with the text generated by OCR to extract medication names, dosage details, and administration frequency. Both modules used fixed configurations throughout all experiments to ensure stable and reproducible evaluations of the integrated system.

Benchmarking Protocol

Computational Performance Metrics

System efficiency was measured using the following metrics:

- Time to extract text: The total time from submitting the prescription image to getting the OCR-generated text.
- End-to-end processing time: The total time needed for image pre-processing, OCR, and clinical NLP entity extraction.
- Processing throughput: The number of prescription images processed per minute under the same experimental conditions.

- Peak CPU memory usage: The maximum resident set size of the backend process during the pipeline's execution.
- Peak system utilization: The average CPU use during OCR and NLP processing.

Diagnostic Baseline

Baseline experiments were carried out with fixed pre-processing steps and consistent model configurations throughout all runs. The OCR and clinical NLP modules ran one after the other without correction to capture the normal behaviour of the pipeline under controlled conditions. These baseline measurements serve as a reference for later analyses of the system.

D. Output Quality Evaluation

Output quality was assessed based on the accuracy and consistency of the extracted prescription information. The quality of the OCR output was judged by recognition error behaviour, while the quality of the NLP output was evaluated using the correctness of medication names, dosage details, and administration frequency. We focused on errors that affect clinically important entities.

E. Qualitative Analysis

A qualitative inspection of system outputs helped to identify common failure modes in the OCR-NLP pipeline. This analysis looked at how handwriting ambiguity, character-level recognition errors, and token boundary issues impacted downstream medical entity extraction.

F. Reproducibility

All experiments were conducted with fixed pre-processing steps, consistent model configurations, and standard evaluation procedures. Using publicly available benchmarks and deterministic execution settings ensures reproducibility and allows for independent verification of the experimental outcomes.

Results and Analysis

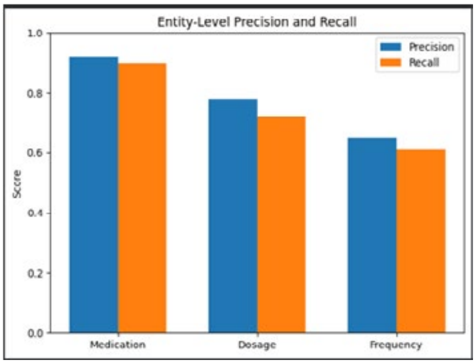

Figure 6.1 Entity Level Precision and Recall

This chart compares the precision and recall for extracting medication names, dosage, and frequency. Medication entities perform best, while dosage and frequency show reduced accuracy due to higher sensitivity to OCR noise.

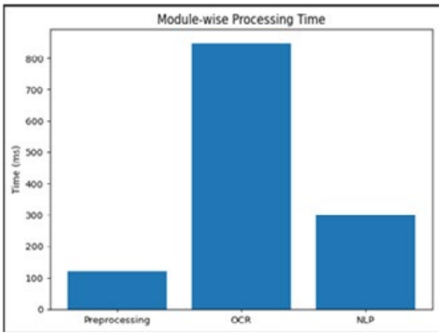

Figure 6.2 Module wise Processing Time

This bar chart displays the average processing time for each pipeline stage. OCR is the computational bottleneck, consuming significantly more time than preprocessing and clinical NLP modules.

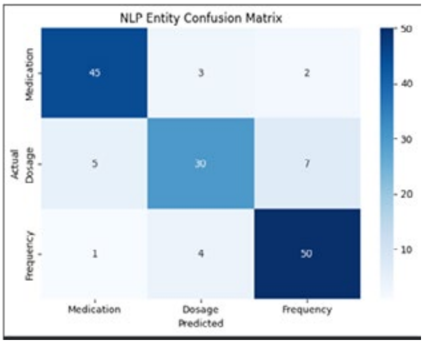

Figure 6.3 NLP Entity Confusion Matrix

The confusion matrix visualizes misclassification patterns in clinical entity extraction. Most errors occur between dosage and frequency entities, indicating that semantic similarity and OCR noise contribute to confusion.

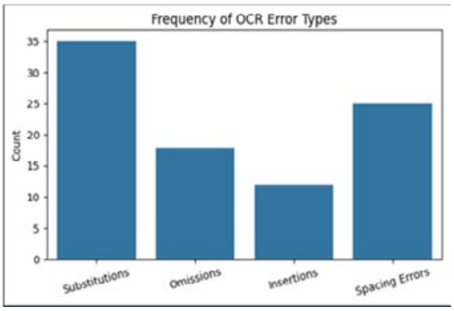

Figure 6.4 Frequency of OCR Error Types

This chart shows the distribution of common OCR error types. Substitutions and spacing errors are the most frequent, often leading to fragmented or misinterpreted medical terms in downstream processing.

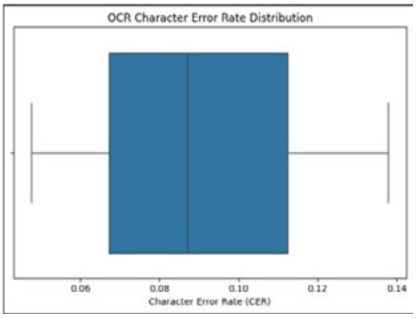

Figure 6.5 OCR Character Error Rate Configuration

This box plot illustrates the variability in Character Error Rate (CER) across different handwritten prescriptions. The spread and presence of outliers highlight how handwriting quality directly influences OCR performance.

B. Entity-Level Behaviour of Clinical NLP Extraction

In clinical NLP evaluation, extraction performance varies greatly across different prescription entity types. Medication names are recognized with higher reliability, thanks to stronger lexical cues and consistent context. However, dosage and administration frequency entities show lower stability because they rely on numeric values, abbreviations, and their position in the prescription text. OCR-related noise has a significant impact on dosage entities, as small recognition errors can drastically change numeric values or unit representations. This sensitivity reveals a limitation of pipeline-based prescription understanding systems since the quality of the NLP components depends on the accuracy of the OCR output. These findings are consistent with what existing clinical NLP literature describes, which notes lower performance for numeric and temporal entities compared to named medical entities.

C. Analysis of OCR-NLP Error Propagation

One major finding from the system-level evaluation is the clear spread of OCR errors into the performance of downstream entity extraction. As recognition errors increase, clinical NLP recall experiences a non-linear decline, especially for dosage and frequency entities. While medication name extraction remains reasonably robust under moderate OCR noise, it sharply declines when recognition errors reach a point where token boundaries and semantic cues become distorted. This behaviour shows that the reliability of a pipeline cannot be assessed solely based on the accuracy of individual components. Instead, the system's strength relies on how well the recognition and interpretation stages work together. The observed amplification of errors underscores the need for integrated evaluation. It also encourages future efforts toward developing error-aware NLP models and post-recognition correction strategies tailored for medical prescriptions.

D. Computational Efficiency and Deployment Feasibility

The end-to-end prescription processing pipeline shows stable computational performance on standard consumer hardware without specialized acceleration. OCR processing takes up most of the total execution time, reflecting the complexity of identifying text in images. Clinical NLP processing adds relatively low overhead, which indicates that text-based analysis can scale efficiently with input size. Memory use stays consistent across experimental runs, suggesting the system can operate in lightweight healthcare settings without needing specialized infrastructure. These results indicate that automated prescription understanding systems could be implemented in small clinics, pharmacies, or patient-facing applications where computational resources are limited.

E. Failure Mode and Qualitative Error Analysis

A qualitative review of the system's outputs highlights recurring failure modes that reveal practical limitations. OCR-related character substitutions and spacing errors often result in fragmented dosage expressions or incomplete frequency indicators. In cases of very cursive handwriting, medication names may only be partially recognized, leading to unclear or incomplete entity extraction. Moreover, domain-specific abbreviations and shorthand commonly found in prescriptions challenge both OCR and NLP components. These insights stress the importance of creating domain-adaptive recognition models and context-aware post-processing techniques. Despite these challenges, the system still successfully extracts meaningful prescription information for most samples, showing its potential as an assistive tool rather than a fully autonomous medical decision system.

Discussion

The experimental results offer important insights into how automated handwritten prescription understanding systems work in practice. Rather than only looking at peak accuracy, the findings emphasize robustness, error sensitivity, and system interactions in real-world healthcare environments.

One key observation is that variability in handwritten prescriptions poses a significant challenge for OCR-based digitization. Even with established recognition methods, performance depends on handwriting style, stroke density, and spatial layout. This shows that recognizing prescriptions differs from understanding general handwritten text. It requires evaluation strategies that account for high variability instead of assuming consistent performance.

The analysis also reveals that the performance of clinical NLP is closely linked to the quality of OCR output. While extracting medication names shows relative robustness in the presence of some recognition noise, details about dosage and administration frequency are much more vulnerable to OCR errors. This sensitivity is critical for patient safety, as numerical and time-related information is often essential for correct medication use. These findings indicate that simply improving downstream NLP is not enough without also enhancing upstream recognition reliability.

A key contribution of this work is the focus on how errors move through the OCR-NLP pipeline. The observed drop in extraction performance indicates that small recognition errors can lead to larger interpretation failures after a certain point. This behaviour stresses the necessity for integrated system evaluation. Checking isolated component metrics might give an overly optimistic view of actual system reliability.

From a deployment standpoint, the system shows practical viability on consumer-grade hardware, making it suitable for resource-limited settings like small clinics, pharmacies, and patient-facing applications. The analysis notes that OCR is the main performance bottleneck. This suggests that future improvements should focus on enhancing recognition efficiency and robustness rather than just downstream text processing.

However, the study does have limitations. The lack of publicly available, end-to-end prescription datasets limits large-scale quantitative evaluation and requires experimentation based on benchmarks. Moreover, the system does not include domain-specific post-processing or context-aware correction techniques, which could help reduce some errors caused by OCR. Nonetheless, these limitations reflect broader challenges in the prescription digitization field rather than issues unique to the proposed framework. Overall, the discussion highlights that automated prescription understanding should be seen as assistive technology that supports human interpretation instead of replacing it. By concentrating on system behaviour, error propagation, and realistic evaluation conditions, this work provides valuable insights for developing safer and more reliable digital healthcare solutions.

Ethical Considerations

Automated analysis of medical prescriptions involves sensitive healthcare information and requires careful ethical consideration. The proposed system focuses on patient privacy, data security, and responsible use of artificial intelligence.

First, no actual patient prescription data was collected, stored, or processed during the experimental evaluation. Due to privacy laws and ethical limits around medical records, the evaluation uses only publicly available benchmark datasets for handwritten text recognition and clinical natural language processing. This approach ensures that no personally identifiable information or protected health information is exposed during development or analysis.

Second, the system is strictly an assistive decision-support tool and does not replace professional medical judgment. Extracted prescription information is presented to improve readability and accessibility, especially for patients who may struggle to interpret handwritten prescriptions. The system does not provide medical advice, alter prescribed medications, or make independent treatment choices. This design lowers the risk of misuse and maintains clinical responsibility with qualified healthcare professionals.

Third, the study acknowledges the potential risks related to recognition and interpretation errors, mainly in dosage and administration frequency extraction. Instead of downplaying these limitations, the evaluation clearly examines how errors can spread across the OCR and NLP pipeline. By identifying vulnerable stages and failure points, the work encourages transparency and supports the development of safer prescription digitization systems.

Additionally, ethical AI principles such as fairness, accountability, and transparency are considered throughout the system's design and evaluation. The modular architecture allows individual components to be checked and evaluated, promoting accountability. Performance limitations are openly reported to avoid overstating system capabilities or encouraging unsafe reliance on automated outputs.

Finally, the analysis of deployment feasibility emphasizes the use of standard hardware and lightweight software components. This supports equal access to digital healthcare technologies and aligns with broader ethical goals of closing healthcare access gaps, especially in resource-limited settings. Overall, the ethical framework in this work prioritizes patient safety, privacy protection, transparency, and responsible AI use. It ensures that automated prescription understanding contributes positively to digital healthcare systems without introducing excessive risk.

Conclusion

This paper presented a complete framework for understanding handwritten prescriptions automatically. It combines Optical Character Recognition with clinical Natural Language Processing in one system. Instead of emphasizing new algorithms, the focus was on how the system behaves, its reliability, and how errors spread, which are crucial for using digital health applications in real life. The analysis shows that recognizing handwritten prescriptions is very sensitive to different writing styles, with OCR errors significantly affecting the extraction of medical information. While extracting medication names is relatively stable, the extraction of dosage and administration frequency is much more affected by recognition errors. These results stress that assessing OCR and NLP separately is not enough. We need to look at them together to truly understand how reliable the system is in practice. From a computing standpoint, the proposed system runs well on standard consumer hardware, making it feasible for use in healthcare settings with limited resources. The findings indicate that future improvements should focus on making recognition more robust and reducing errors, rather than just refining the text processing that comes after. Overall, this work gives a realistic view of automated prescription understanding systems by emphasizing how different parts work together, what can go wrong, and important ethical considerations. By using benchmark-guided evaluation and clearly reporting limitations, this study offers insights that can help create safer, more reliable, and patient-focused digital health solutions. Future research could look into recognition models that adapt to different domains, correction methods after recognition, and controlled clinical trials to further enhance system reliability and usability.

References

1. Datta, A., Hasan, M. M., Mahmud, N., Ferdosi, B. J., Islam, R., & Mehedi, S. T. (2025). Preserving medical information from doctor's prescription ensuring relation among the terminology. Computers in Biology and Medicine, 188, Art. no. 109812.
2. Wang, Y., Steinhubl, S. R., Defilippi, C., Ng, K., Ebadollahi, S., Stewart, W. F., & Byrd, R. J. (2015). Prescription extraction from clinical notes: toward automating EMR medication reconciliation. In Proc. AMIA Summit Translational Science (pp. 188–193).
3. Xu, H., Stenner, S. P., Doan, S., Johnson, K. B., Waitman, L. R., & Denny, J. C. (2010). MedEx: a medication information extraction system for clinical narratives. Journal of the American Medical Informatics Association, 17(1), 19–24.
4. Hsu, E., Malagaris, I., Kuo, Y.-F., Sultana, R., & Roberts, K. (2022). Deep learning-based NLP data pipeline for EHR-scanned document information extraction. JAMIA Open, 5(2), ooac045.
5. Haaker, T. S., Choi, J. S., Nanjo, C. J., Warner, P. B., Abu-Hanna, A., & Kawamoto, K. (2025). Approaches for extracting daily dosage from free-text prescription signatures in heart failure with reduced ejection fraction: a comparative study. JAMIA Open, 8(1), ooae153.
6. Tabassum, S., Abedin, N., Rahman, M. M., Rahman, M. M., Ahmed, M. T., Islam, R., & Ahmed, A. (2022). An online cursive handwritten medical words recognition system for busy doctors in developing countries for ensuring efficient healthcare service delivery. Scientific Reports, 12, Art. no. 3601.
7. Kim, G., Hong, T., Yim, M., Nam, J., Park, J., Yim, J., Hwang, W., Yun, S., Han, D., & Park, S. (2022). OCR-free Document Understanding Transformer. In Proc. European Conference on Computer Vision (ECCV) (pp. 106–122).
8. Abujaber, A. A., & Nashwan, A. J. (2024). Ethical framework for artificial intelligence in healthcare research: a path to integrity. World Journal of Methodology, 14(3), Art. no. 94071.
9. Abràmoff, M. D., Tarver, M. E., Loyo-Berrios, N., Trujillo, S., Char, D., Obermeyer, Z., Eydelman, M. B., & Maisel, W. H. (2023). Considerations for addressing bias in artificial intelligence for health equity. NPJ Digital Medicine, 6, Art. no. 170.
10. Dankwa-Mullan, I. (2024). Health equity and ethical considerations in using artificial intelligence in public health and medicine. Preventing Chronic Disease, 21, 240245.