iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL
FOR RESEARCH
IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** VI **Month of publication:** June 2026

DOI: https://doi.org/10.22214/ijraset.2026.83831

PACCF: Prescription-Aware Adaptive Confidence Fusion for Robust Medical Prescription OCR

Manish Joshi
Dept. of Computer Science and Engineering, Uttarakhand Technical University, Dehradun, India

Abstract: The digitization of handwritten medical prescriptions remains a challenging task due to poor handwriting, variations in writing styles, image quality issues, and the presence of specialized pharmaceutical terminology. Although recent Optical Character Recognition (OCR) technologies have achieved significant progress in document understanding, their performance on medical prescriptions is often limited because most systems rely on a single recognition model and are therefore vulnerable to model-specific errors. To address these challenges, this paper proposes a novel Prescription-Aware Adaptive Confidence Fusion (PACCF) framework for accurate and reliable medical prescription text recognition.

The proposed approach combines the strengths of multiple OCR engines, namely TrOCR, PARSeq, EasyOCR, and PaddleOCR, to generate diverse recognition hypotheses from prescription text regions. Instead of selecting predictions based solely on confidence scores, PACCF introduces an adaptive fusion strategy that evaluates each candidate using four complementary reliability measures: confidence reliability, inter-model agreement, Gaussian consensus reliability, and medical lexical similarity. These factors are dynamically weighted through a Softmax-based mechanism, enabling data-driven decision making without manually tuned parameters. To further enhance recognition accuracy, a BK-Tree-based medical validation module is employed for efficient approximate matching and correction of pharmaceutical terms.

Experimental results on a challenging handwritten prescription dataset demonstrate that the proposed framework significantly improves both character-level and word-level recognition performance compared with individual OCR models and conventional fusion techniques. The findings suggest that integrating ensemble diversity, adaptive confidence modeling, and domain-specific medical knowledge provides an effective and scalable solution for automated prescription digitization and healthcare document processing.

Keywords: Medical Prescription Recognition, Optical Character Recognition (OCR), Multi-OCR Ensemble, Adaptive Confidence Fusion, PACCF, Medical Text Recognition, BK-Tree Validation, Healthcare Document Analysis, Pharmaceutical Term Recognition, Artificial Intelligence in Healthcare.

I. INTRODUCTION

A. Problem Statement

The digitization of medical prescriptions plays a critical role in healthcare information management, pharmacy automation, clinical decision support systems, and electronic health record generation. Despite significant advances in Optical Character Recognition (OCR), reliable extraction of prescription text remains a challenging task due to handwritten content, image degradation, inconsistent writing styles, low contrast, and the presence of specialized pharmaceutical terminology. Errors in prescription recognition can lead to incorrect medicine identification, affecting the reliability of downstream healthcare applications. Therefore, developing robust and accurate prescription text recognition systems remains an important research problem.

B. Research Gap

Recent deep learning-based OCR models, including transformer-based architectures, have demonstrated substantial improvements in document text recognition.

However, existing approaches predominantly rely on a single OCR engine, making them vulnerable to model-specific recognition failures under varying prescription image conditions. Although OCR ensemble techniques have been explored in related domains, most fusion strategies employ static weighting schemes or confidence-based voting mechanisms that do not adequately consider inter-model consensus, statistical confidence reliability, and domain-specific medical knowledge. Furthermore, existing prescription OCR systems rarely incorporate scalable medical validation mechanisms capable of efficiently handling large pharmaceutical vocabularies.

C. Proposed Solution

To address these limitations, this paper proposes a Prescription-Aware Adaptive Confidence Fusion (PACCF) framework for robust medical prescription text recognition. The proposed framework employs a heterogeneous OCR ensemble comprising TrOCR, PARSeq, EasyOCR, and PaddleOCR to generate complementary recognition hypotheses. A novel adaptive fusion mechanism evaluates each candidate prediction using four reliability indicators: confidence reliability, hybrid OCR agreement, Gaussian consensus confidence estimation, and medical lexical similarity. Unlike conventional fusion methods that depend on manually selected weighting coefficients, PACCF utilizes adaptive Softmax-based factor weighting to dynamically determine the contribution of each reliability factor. In addition, a BK-Tree accelerated medical validation module is incorporated to perform efficient approximate retrieval and correction of pharmaceutical terms.

D. Contributions

The major contributions of this work are summarized as follows:

1) A heterogeneous multi-OCR ensemble architecture integrating transformer-based and sequence-based OCR models for prescription text recognition.
2) The proposed Prescription-Aware Adaptive Confidence Fusion (PACCF) algorithm that combines confidence reliability, hybrid agreement analysis, Gaussian consensus modeling, and medical lexical similarity for adaptive decision fusion.
3) An adaptive Softmax factor-weighting mechanism that eliminates manually tuned fusion coefficients and enables data-driven prediction selection.
4) A BK-Tree based medical validation framework for scalable approximate retrieval and correction of medicine names.
5) A comprehensive experimental evaluation and ablation analysis demonstrating the effectiveness of individual PACCF components and the overall framework for medical prescription recognition.
6) **Novelty Statement:** The novelty of the proposed PACCF framework lies in its unified integration of adaptive factor-wise Softmax fusion, Gaussian consensus reliability estimation, hybrid inter-model agreement analysis, and BK-Tree-based medical validation for handwritten prescription OCR. Unlike existing confidence-based ensemble methods, PACCF dynamically balances multiple reliability indicators while incorporating domain-specific pharmaceutical knowledge during the fusion process, enabling more accurate and robust prescription recognition.

II. RELATED WORK

A. Deep Learning OCR

Recent advances in deep learning have significantly transformed Optical Character Recognition (OCR) systems. Traditional OCR approaches relied on handcrafted feature extraction, character segmentation, and statistical classifiers, which often struggled with complex document layouts and degraded images. The introduction of Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and Transformer-based architectures has substantially improved recognition accuracy across various document analysis tasks. Transformer-based OCR models such as TrOCR employ a Vision Transformer encoder and Transformer decoder architecture to perform end-to-end text recognition without requiring explicit character segmentation. Similarly, PARSeq introduces a permuted autoregressive sequence modeling strategy that improves robustness for irregular and challenging text sequences. Other widely adopted OCR frameworks, including EasyOCR and PaddleOCR, provide practical and computationally efficient recognition capabilities for real-world applications. Despite these advances, individual OCR models continue to exhibit performance variations depending on handwriting style, image quality, and document complexity.

B. OCR Ensemble Methods

To overcome the limitations of individual OCR systems, several studies have explored ensemble-based recognition strategies. Ensemble methods combine predictions from multiple OCR engines to exploit complementary model strengths and improve overall recognition reliability. Common approaches include majority voting, confidence-based selection, weighted voting, and rank aggregation techniques. Although ensemble OCR systems generally outperform standalone models, most existing methods rely on fixed weighting schemes or simple confidence aggregation mechanisms. Such approaches assume static model reliability and often fail to account for prediction consensus, uncertainty estimation, and contextual relevance. Consequently, ensemble decisions may remain vulnerable to highly confident yet incorrect predictions. These limitations motivate the development of adaptive fusion mechanisms capable of dynamically evaluating prediction reliability using multiple complementary indicators.

C. Medical Prescription Recognition

Medical prescription recognition represents a specialized application of document analysis characterized by unique challenges including handwritten text, image artifacts, abbreviated medical terminology, and complex pharmaceutical nomenclature. Accurate prescription digitization is essential for pharmacy automation, electronic health records, medication management systems, and healthcare analytics.

Previous research in prescription recognition has primarily focused on image enhancement techniques, handwriting recognition models, and domain-specific OCR solutions. Several systems incorporate post-processing methods such as dictionary matching and spell correction to improve medicine name recognition. However, most existing approaches depend on a single OCR engine and perform domain validation only after recognition. As a result, recognition errors generated during earlier stages frequently propagate throughout the pipeline, reducing overall system reliability.

D. Limitations of Existing Approaches

Despite substantial progress in OCR and document intelligence research, several limitations remain unresolved in current prescription recognition systems.

First, most OCR-based prescription recognition frameworks rely on a single recognition model, making them susceptible to model-specific failure modes under varying image conditions. Second, existing ensemble methods commonly employ static weighting strategies that do not adapt to prediction uncertainty or inter-model agreement. Third, confidence scores generated by OCR models are often treated as direct indicators of correctness, despite their susceptibility to overconfidence and calibration issues. Fourth, domain-specific medical knowledge is typically utilized only during post-processing, preventing it from influencing the fusion decision itself.

Finally, many dictionary-based validation approaches rely on linear search mechanisms that become computationally inefficient when applied to large pharmaceutical databases.

To address these challenges, the proposed Prescription-Aware Adaptive Confidence Fusion (PACCF) framework integrates heterogeneous OCR recognition, adaptive confidence modeling, consensus-based reliability estimation, medical lexical similarity analysis, and BK-Tree accelerated validation into a unified decision-making framework for robust prescription text recognition.

III. PACCF FRAMEWORK

A. System Architecture

The overall architecture of the proposed Prescription-Aware Adaptive Confidence Fusion (PACCF) framework is illustrated in Fig. 1. The framework follows a multi-stage pipeline designed for robust medical prescription text recognition. Initially, the input prescription image undergoes image preprocessing to enhance text visibility and reduce noise artifacts.

The processed image is then passed to the text detection module, which localizes candidate text regions for recognition. Each detected region is subsequently processed by a heterogeneous Multi-OCR Ensemble comprising TrOCR, PARSeq, EasyOCR, and PaddleOCR, generating multiple candidate predictions along with their corresponding confidence scores. These predictions are aggregated within the PACCF adaptive fusion module, where four complementary reliability factors—Confidence Reliability, Hybrid Agreement, Gaussian Reliability, and Medical Similarity—are computed to assess prediction quality. An adaptive Softmax weighting mechanism dynamically determines the contribution of each factor and produces a unified fusion score for every candidate.

The highest-ranked prediction is then validated through a BK-Tree-based medical retrieval module that performs efficient approximate matching against a medicine knowledge base. Finally, the validated prediction is returned as the corrected prescription text. This architecture integrates recognition diversity, adaptive confidence modeling, statistical reliability estimation, and domain-specific medical validation into a unified framework for accurate prescription digitization.

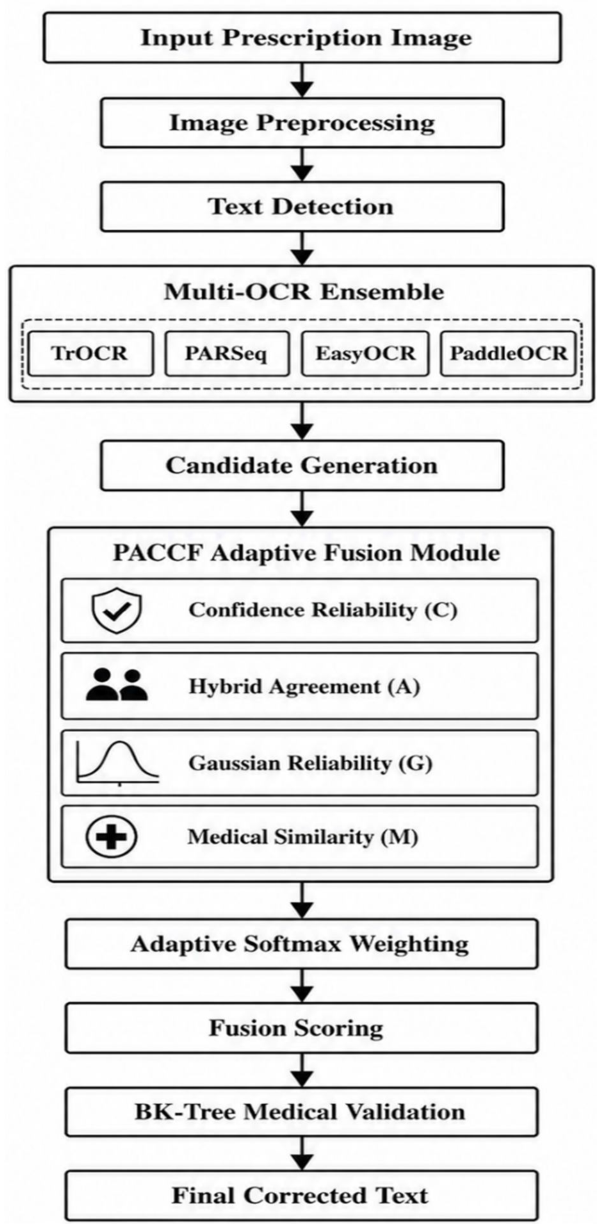

Figure 1. Overall architecture of the proposed Prescription-Aware Adaptive Confidence Fusion (PACCF) framework for medical prescription text recognition.

B. Multi-OCR Candidate Generation

To improve recognition robustness and mitigate model-specific errors, the proposed framework employs a heterogeneous Multi-OCR Ensemble consisting of TrOCR, PARSeq, EasyOCR, and PaddleOCR. Rather than relying on a single recognition model, each detected text region is independently processed by all OCR engines to generate multiple candidate predictions. This strategy exploits the complementary strengths of transformer-based, sequence-based, and industrial OCR architectures, thereby increasing the likelihood of obtaining a correct prediction under diverse prescription image conditions.

TrOCR is utilized as a transformer-based recognition model that leverages a Vision Transformer encoder and Transformer decoder architecture for handwritten and printed text recognition. PARSeq employs a permutation-based autoregressive sequence modeling approach that provides strong sequence-level recognition performance. EasyOCR contributes a CRNN-based recognition pipeline capable of handling noisy and low-quality text regions, while PaddleOCR provides a lightweight industrial-grade OCR solution optimized for practical document analysis tasks.

For each detected text region R_n, the ensemble generates a set of candidate predictions:

$$Y_R = \{(y_i, c_i)\}_{i=1}^{K}$$

where y_i represents the text prediction produced by the i^{th} OCR engine and c_i denotes the corresponding confidence score. The resulting candidate set contains multiple recognition hypotheses that may differ in textual content and confidence levels. These hypotheses are subsequently forwarded to the PACCF module, where adaptive confidence fusion is performed to identify the most reliable prediction.

The use of multiple OCR engines enables the framework to capture complementary recognition patterns and reduces dependence on the failure characteristics of any individual model, thereby providing a more reliable foundation for downstream fusion and medical validation stages.

C. Prescription-Aware Adaptive Confidence Fusion (PACCF)

The Prescription-Aware Adaptive Confidence Fusion (PACCF) module constitutes the core decision-making component of the proposed framework. Given a set of candidate predictions generated by the Multi-OCR Ensemble, PACCF evaluates the reliability of each candidate using multiple complementary factors and dynamically determines the most probable recognition output.

Unlike conventional ensemble OCR systems that rely solely on confidence-based voting or manually assigned weights, PACCF integrates statistical confidence modeling, inter-model consensus analysis, and domain-aware medical validation within a unified adaptive fusion framework. The objective is to select the candidate that maximizes both recognition reliability and medical plausibility.

The PACCF architecture is illustrated in Fig. 2.

1) Confidence Reliability Factor

The confidence reliability factor quantifies the certainty associated with a prediction generated by an OCR engine. Let c_i denote the confidence score assigned to candidate y_i . The confidence reliability is defined as

$C_i = c_i$

where larger values indicate higher prediction certainty. Although confidence provides useful information regarding recognition quality, it is insufficient as a standalone decision criterion because OCR systems may occasionally produce highly confident but incorrect predictions.

2) Hybrid Agreement Factor

To estimate consensus among OCR models, PACCF computes a hybrid agreement score that combines character-level edit similarity and structural sequence similarity.

For two candidate predictions y_i and y_i normalized Levenshtein similarity is computed as

$$L_{ij} = 1 - \frac{d_{edit}(y_i, y_j)}{max(| y_i |, | y_j |)}$$

Character N-Gram similarity is computed using Jaccard similarity over character n-gram sets:

$$N_{ij} = \frac{|G(y_i) \cap G(y_j)|}{|G(y_i) \cup G(y_j)|}$$

(Where $G(y)$denotes the set of n-grams extracted from string y.)

The pairwise agreement score is then defined as

$$A_{\{ij\}} = 0.7L_{\{ij\}} + 0.3N_{\{ij\}}$$

The overall agreement factor for candidate y_i is obtained by averaging agreement scores across all remaining OCR candidates. This factor promotes predictions that are consistently supported by multiple recognition engines.

3) Gaussian Consensus Reliability

Confidence scores generated by OCR models may vary significantly across recognition engines. PACCF therefore introduces a Gaussian consensus reliability factor that evaluates how closely a candidate confidence aligns with the ensemble consensus.

Let

$$\mu_c = \left(\frac{1}{K}\right)\sum_{i=1}^{K} c_i$$

represent the mean confidence and

$$\sigma_c^2 = \left(\frac{1}{K}\right)\sum_{i=1}^{K} \quad (c_i - \mu_c)^2$$

represent the confidence variance.
The Gaussian reliability score is computed as

$$G_i = exp\ exp \left(-\frac{(c_i - \mu_c)^2}{2\sigma_c^2 + \varepsilon}\right)$$

Where $\varepsilon = 10^{-8}$ is a numerical stabilization constant.
Candidates whose confidence values closely follow the ensemble consensus receive higher reliability scores, whereas outlier predictions are penalized.

4) Medical Lexical Similarity

Medical prescriptions contain domain-specific terminology that is often absent from generic OCR vocabularies. To incorporate domain knowledge into the fusion process, PACCF evaluates the lexical similarity between OCR predictions and entries stored within a medicine knowledge repository.
The medical similarity factor combines four complementary similarity measures:
- Fuzzy string similarity
- Edit-distance similarity
- Database matching similarity
- Character N-Gram similarity

The final medical similarity score is calculated as

$$M_i = \left(S_fuzzy + S_edit + S_db + S_ngram\right)/4$$

This factor enables medically plausible predictions to receive additional support during the fusion process.

5) Adaptive Softmax Weighting

Rather than employing manually tuned weighting coefficients, PACCF dynamically determines the importance of each reliability factor through an adaptive Softmax weighting mechanism. For each candidate prediction, the reliability factors—Confidence Reliability (C_i), Hybrid Agreement (A_i), Gaussian Reliability (G_i), and Medical Similarity (M_i)—are transformed into normalized adaptive weights using the Softmax function to obtain their corresponding contributions to the final fusion score:

$$W_C = \frac{e^{C_i}}{e^{C_i} + e^{A_i} + e^{G_i} + e^{M_i}}$$

$$W_A = \frac{e^{A_i}}{e^{C_i} + e^{A_i} + e^{G_i} + e^{M_i}} W_G = \frac{e^{G_i}}{e^{C_i} + e^{A_i} + e^{G_i} + e^{M_i}} W_M = \frac{e^{M_i}}{e^{C_i} + e^{A_i} + e^{G_i} + e^{M_i}}$$

The resulting weights satisfy

$$W_C + W_A + W_G + W_M = 1$$

thereby ensuring normalized and data-driven contribution allocation. This adaptive weighting mechanism enables PACCF to automatically emphasize the most informative reliability factors for each candidate prediction while reducing dependence on manually selected weighting coefficients.
where W_C, W_A, W_G, and W_M represent the normalized adaptive weights assigned to the Confidence Reliability, Hybrid Agreement, Gaussian Reliability, and Medical Similarity factors, respectively. The weights are computed through a Softmax transformation of the corresponding reliability scores, allowing PACCF to dynamically balance statistical confidence, inter-model consensus, and domain-specific medical knowledge during the fusion process.
Unlike conventional ensemble methods that apply Softmax normalization across competing candidate predictions, PACCF applies Softmax normalization across the reliability factors associated with each candidate prediction. This design enables the framework to adaptively determine the relative importance of Confidence Reliability, Hybrid Agreement, Gaussian Reliability, and Medical Similarity for each candidate independently. Consequently, the final fusion score reflects not only the magnitude of individual reliability measures but also their relative contribution to the overall prediction reliability, resulting in a more balanced and domain-aware decision-making process.

6) Final Fusion Score

The final PACCF decision score is obtained by integrating all reliability factors using their adaptive Softmax weights:

$$F_i = W_C C_i + W_A A_i + W_G G_i + W_M M_i$$

The candidate with the highest fusion score is selected as the final ensemble prediction and subsequently forwarded to the medical validation stage.

By jointly modeling confidence reliability, inter-model agreement, statistical consensus, and medical lexical similarity, PACCF provides a robust and domain-aware decision mechanism for medical prescription text recognition.

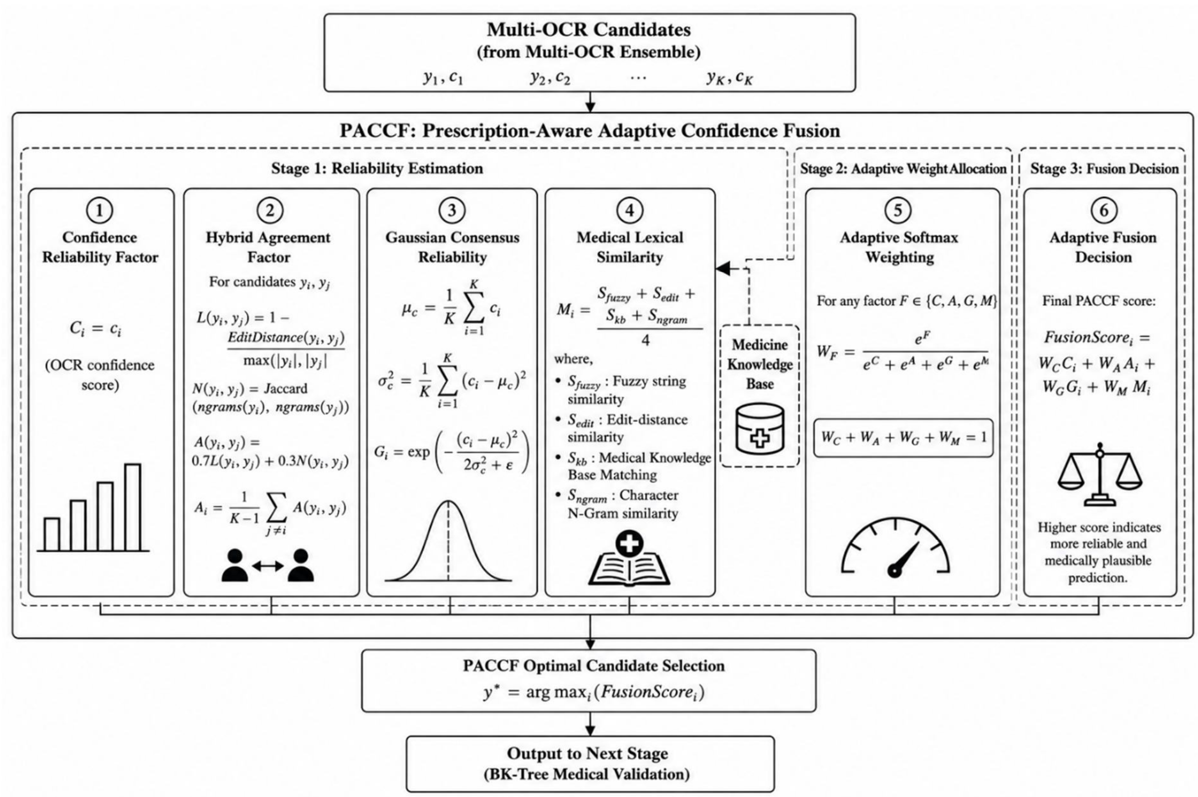

Figure 2. PACCF module for adaptive fusion of multi-OCR predictions.

D. BK-Tree Medical Validation

Following the PACCF fusion stage, the selected candidate prediction undergoes domain-specific validation using a BK Tree-based medical retrieval framework. The objective of this stage is to identify and correct residual recognition errors that may persist after adaptive fusion, particularly for medicine names containing visually similar characters or uncommon pharmaceutical terminology.

A Burkhard-Keller Tree (BK-Tree) is employed as an efficient indexing structure for approximate string matching within a medicine knowledge repository. Each medicine entry is stored as a node in the BK-Tree, where edges are organized according to edit-distance relationships.

During validation, the fused OCR prediction is queried against the tree using Levenshtein distance as the similarity metric. Instead of performing an exhaustive comparison with every medicine entry, the BK-Tree selectively traverses only relevant branches that satisfy the specified edit-distance threshold. Let (q) denote the fused OCR prediction and (D) represent the medicine database. The retrieval process identifies a set of candidate medicines:

$$M(q) = \{m \in D \mid d_{edit}(q, m) \leq \tau\}$$

where τ denotes the maximum allowable edit distance.

For each retrieved candidate, a medical validation score is computed using lexical similarity measures. The medicine with the highest validation score is selected as the final correction candidate. Correction is applied only when both the lexical similarity score and edit-distance constraints satisfy predefined acceptance thresholds, thereby preventing unnecessary modifications to correctly recognized text.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VI June 2026- Available at www.ijraset.com

Compared with conventional linear dictionary search methods having complexity (O(D)), the proposed BK-Tree retrieval mechanism significantly reduces search overhead and improves scalability for large pharmaceutical vocabularies. Consequently, the framework achieves efficient domain-aware validation while maintaining practical computational performance for real-world prescription digitization systems.

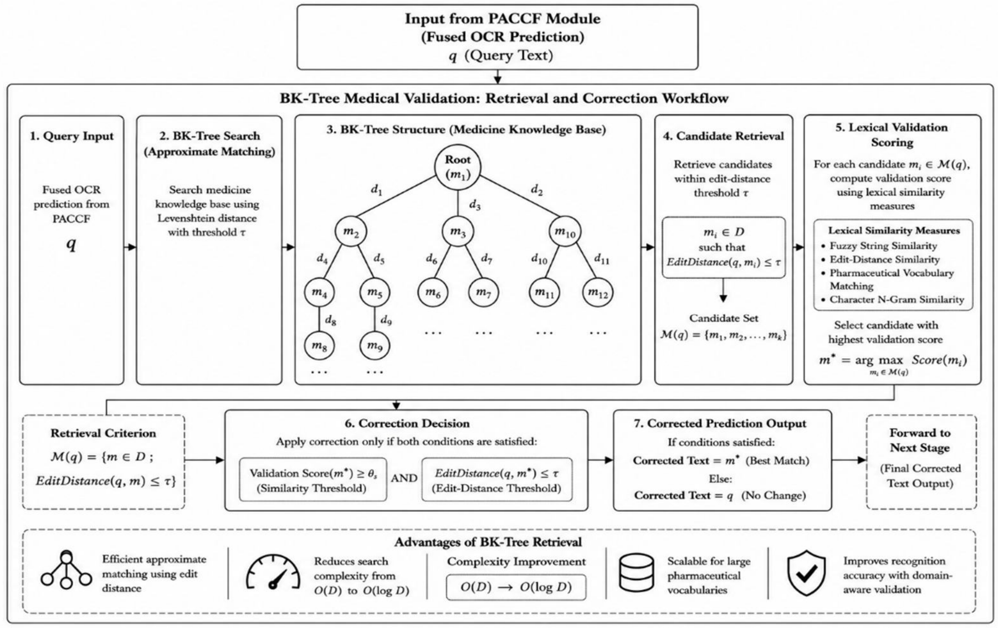

Figure 3. BK-Tree-based correction workflow for validating fused OCR predictions against a medical knowledge repository

IV. COMPUTATIONAL ANALYSIS

This section analyzes the computational characteristics of the proposed Prescription-Aware Adaptive Confidence
Fusion (PACCF) framework. The analysis focuses on the time complexity, space complexity, and scalability of the framework with respect to image size,
number of OCR models, and medicine database size.

A. Time Complexity

The overall computational cost of the proposed framework consists of four major components: image preprocessing, text detection, multi-OCR recognition, and medical validation.

Let:

- (W \times H) denote the image resolution.
- (N) denote the number of detected text regions.
- (K) denote the number of OCR models.
- (D) denote the number of medicine entries in the knowledge base.

The preprocessing and text detection stages operate on image pixels and require

$$O(WH)$$

time complexity.

For each detected text region, the framework executes (K) OCR engines. Let T_{OCR} represent the average recognition cost of a single OCR model. The recognition complexity becomes

$$O(NKT_{OCR})$$

For each candidate prediction, PACCF computes confidence reliability, agreement, Gaussian reliability, and medical lexical similarity. Let Kdenote the number of OCR models in the ensemble. The fusion stage has complexity $O(NK^2)$. For the current implementation, $K = 4$.

The BK-Tree retrieval stage performs approximate medicine matching with average-case complexity

$$O(logD)$$

per query. Consequently, the medical validation stage requires

$$O(NK \log D)$$

time.

Therefore, the total computational complexity of the proposed PACCF framework is

$$O(WH + NKT_{OCR} + NK^2 + NK \log \log D)$$

Since (K) is a small constant, the dominant computational cost originates from OCR inference, while BK-Tree retrieval introduces only logarithmic search overhead.

B. Space Complexity

The space complexity of the framework is determined by image storage, OCR candidate generation, and medicine knowledge indexing.

The input image requires

O(WH)

memory.

The OCR ensemble generates (K) candidate predictions for each detected text region, resulting in

$$O(NK)$$

storage complexity.

The medicine repository containing (D) medicine entries is stored within the BK-Tree structure. Consequently, the medical knowledge storage requirement is

$$O(D)$$

The total space complexity is therefore given by

$$O(WH + NK + D)$$

which scales linearly with the size of the medicine database and the number of generated OCR candidates.

C. Scalability Analysis

Scalability is a critical requirement for real-world prescription digitization systems. Traditional dictionary-based validation methods perform exhaustive searches over all medicine entries, resulting in

$$O(D)$$

retrieval complexity for each query.

In contrast, the proposed framework utilizes BK-Tree indexing, which reduces average retrieval complexity to

$$O(\log \log D)$$

through edit-distance guided traversal.

As the pharmaceutical vocabulary grows, the computational benefit becomes increasingly significant. The proposed retrieval strategy therefore enables efficient deployment on large-scale medicine repositories while maintaining low query latency.

Furthermore, the modular architecture of PACCF allows additional OCR models to be incorporated into the ensemble without modifying the fusion mechanism. Since fusion complexity depends primarily on the number of candidate predictions rather than database size, the framework remains computationally practical for large-scale healthcare document processing environments.

The complexity analysis demonstrates that PACCF achieves improved scalability through logarithmic-time medical retrieval while preserving the recognition benefits of a heterogeneous OCR ensemble.

Method	Validation Strategy	Retrieval Complexity
Dictionary Search	Linear Scan	$O(D)$
Hash Lookup	Exact Match Only	$O(1)$
Trie Search	Prefix Matching	$O(L)$
Proposed PACCF	BK-Tree Approximate Retrieval	$O(\log D)$

where D is the medicine database size and L is the query length.

Table I. Complexity Comparison

V. EXPERIMENTAL SETUP

This section describes the dataset, evaluation metrics, and baseline OCR models used to assess the effectiveness of the proposed Prescription-Aware Adaptive Confidence Fusion (PACCF) framework.

A. Dataset Description

A custom prescription dataset was constructed for experimental evaluation. The dataset consists of 137 camera-captured prescription images containing handwritten medicine names written primarily using blue ink on plain paper. The images were acquired under diverse real-world imaging conditions, including variations in illumination intensity, viewing angle, background noise, and image quality.

The collected samples exhibit substantial handwriting variability, ranging from moderately legible writing to highly challenging prescription-style handwriting. Such variability introduces significant recognition complexity due to character deformation, inconsistent spacing, stroke discontinuities, and ambiguous letter formations.

Unlike benchmark document datasets captured under controlled conditions, the proposed dataset reflects practical prescription acquisition scenarios commonly encountered in healthcare environments. Consequently, it provides a suitable testbed for evaluating the robustness of OCR systems and adaptive fusion strategies under realistic operating conditions.

All images were manually verified and annotated to establish ground-truth labels for quantitative performance evaluation.

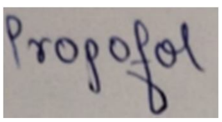

Sample 1
(Easy)

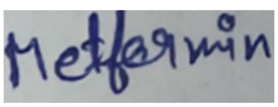

Sample 2
(Medium)

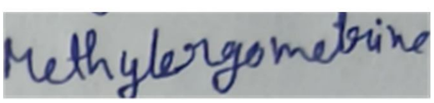

Sample 3
(Difficult)

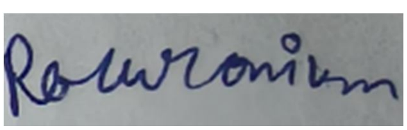

Sample 4
(Poor Writing)

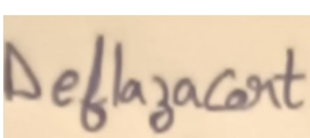

Sample 5
(Blur)

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VI June 2026- Available at www.ijraset.com

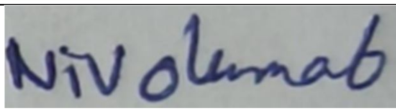

Sample 6
(Severe Handwriting)

Table II. Representative samples from the prescription dataset.

2) Word Recognition Accuracy (WRA)

Word Recognition Accuracy evaluates the proportion of completely correct word predictions.

$$WRA = \frac{W_{correct}}{W_{total}} \times 100$$

where (W_{correct}) denotes correctly recognized words and (W_{total}) denotes the total number of evaluated words.

3) Precision

Precision measures the proportion of correctly recognized predictions among all predicted outputs.

$$Precision = \frac{TP}{TP + FP}$$

where (TP) and (FP) denote true positives and false positives, respectively.

4) Recall

Recall measures the proportion of correctly recognized words among all ground-truth words.

$$Recall = \frac{TP}{TP + FN}$$

where (FN) denotes false negatives.

5) F1 Score

The F1 Score provides a balanced evaluation by combining Precision and Recall.

$$F1 = \frac{2 \times Precision \times Recall}{Precision + Recall}$$

A higher F1 Score indicates better overall recognition performance.

B. Baseline Methods

To evaluate the effectiveness of PACCF, the proposed framework was compared against four widely used OCR systems operating independently without adaptive fusion or medical validation.

1) TrOCR

TrOCR is a transformer-based OCR model employing a Vision Transformer encoder and Transformer decoder architecture. It serves as a strong baseline for handwritten and printed text recognition.

2) PARSeq

PARSeq is a transformer-based scene and document text recognition model that utilizes permutation-based autoregressive sequence learning for robust character prediction.

3) EasyOCR

EasyOCR is a CRNN-based OCR framework that combines convolutional feature extraction and recurrent sequence modeling. It is widely used due to its flexibility and lightweight deployment requirements.

4) PaddleOCR

PaddleOCR is an industrial-grade OCR framework optimized for practical document analysis applications. It provides efficient text detection and recognition capabilities across diverse document types.

The proposed PACCF framework integrates predictions from all four OCR engines and applies adaptive confidence fusion together with domain-aware medical validation. Comparative evaluation against these baseline models enables quantitative assessment of the contribution of PACCF beyond individual OCR performance.

Parameter	Value
OCR Models	TrOCR, PARSeq, EasyOCR, PaddleOCR
Number of OCR Engines ((K))	4
Fusion Method	PACCF
Validation Method	BK-Tree Retrieval
Similarity Metrics	Fuzzy, Edit, N-Gram, Vocabulary Matching
Evaluation Metrics	CRA, WRA, Precision, Recall, F1

Table III. Experimental Configuration.

C. *Hardware and Software Configuration*

Experiments were conducted on a workstation equipped with an NVIDIA GeForce RTX 3050 GPU (6 GB GDDR6) and 16 GB RAM. The proposed PACCF framework was implemented using Python 3.11 and executed under a CUDA-enabled environment for GPU acceleration. The OCR ensemble consisted of TrOCR, PARSeq, EasyOCR, and PaddleOCR (PP-OCRv5). Deep learning models were implemented using PyTorch and the Hugging Face Transformers library, while image processing operations were performed using OpenCV and NumPy. Medical validation and similarity computation utilized TheFuzz and python-Levenshtein libraries. PaddleOCR was configured with compatibility optimizations for stable execution.

GPU acceleration was employed for TrOCR, PARSeq, and EasyOCR inference, while batch processing was utilized to improve computational efficiency during large-scale evaluation.

VI. RESULTS AND ANALYSIS

This section evaluates the effectiveness of the proposed Prescription-Aware Adaptive Confidence Fusion (PACCF) framework through comparative experiments, component-wise ablation studies, retrieval analysis, and qualitative recognition examples.

A. *Baseline Comparison*

The quantitative results presented in Table IV demonstrate the superiority of the proposed PACCF framework over individual OCR engines across all evaluation metrics. Among the standalone OCR models, PARSeq achieved the highest Character Recognition Accuracy (CRA) of 82.21%, while PaddleOCR obtained the highest Word Recognition Accuracy (WRA) of 31.11%. EasyOCR exhibited the weakest performance due to its limited robustness against highly degraded handwriting, achieving a CRA of only 41.24% and a WRA of 0.74%. Although TrOCR achieved relatively high recall (76.04%), its overall word-level recognition performance remained limited, indicating difficulties in accurately recognizing challenging prescription terms.

In contrast, the proposed PACCF framework achieved a CRA of 93.86% and a WRA of 78.52%, substantially outperforming all baseline methods. Compared with the strongest individual OCR engine (PARSeq), PACCF improved character-level accuracy by 11.65 percentage points and increased word-level accuracy by 52.59 percentage points. Furthermore, PACCF achieved the highest Precision (91.14%), Recall (95.97%), and F1 Score (93.49%), demonstrating superior recognition reliability and consistency.

The significant improvement can be attributed to three key factors. First, the heterogeneous OCR ensemble generates complementary recognition hypotheses, reducing model-specific failure modes. Second, the proposed Prescription-Aware Adaptive Confidence Fusion (PACCF) mechanism effectively combines confidence reliability, inter-model agreement, Gaussian consensus estimation, and medical lexical similarity to identify the most reliable prediction. Third, the BK-Tree-based medical validation stage further refines pharmaceutical terminology through efficient domain-aware correction. These results confirm that adaptive confidence fusion and medical-domain validation substantially enhance prescription text recognition performance under challenging real-world handwriting conditions.

Method	CRA (%)	WRA (%)	Precision (%)	Recall (%)	F1 Score (%)
EasyOCR	41.24	0.74	60.85	41.64	49.45
PaddleOCR	80.47	31.11	89.9	80.37	84.87
PARSeq	82.21	25.93	83.82	83.13	83.48
TrOCR	65.99	11.11	66.04	76.04	70.69
PACCF	93.86	78.52	91.14	95.97	93.49

Table IV. Performance Comparison with Baseline OCR Model.

B. PACCF Ablation Study

The ablation results presented in Table V demonstrate the incremental contribution of each PACCF component toward improving prescription recognition performance. Starting from the Confidence Reliability baseline, the framework achieved a Character Recognition Accuracy (CRA) of 68.15% and a Word Recognition Accuracy (WRA) of 22.22%. Incorporating the Hybrid Agreement factor increased CRA to 82.92% and WRA to 32.59%, indicating that inter-model consensus provides valuable information for identifying reliable OCR predictions.

The addition of the Gaussian Reliability component resulted in a CRA of 85.12% and a WRA of 26.67%. Although the immediate improvement was limited, Gaussian consensus modeling contributed to stabilizing confidence estimation and reducing the influence of confidence outliers within the ensemble. This factor becomes more beneficial when combined with subsequent domain-aware components.

When Medical Lexical Similarity was introduced, the framework achieved a substantial performance improvement, increasing CRA to 90.37% and WRA to 44.44%. This result highlights the importance of incorporating domain-specific medical knowledge into the fusion process, enabling the framework to favor medically plausible predictions even when OCR outputs exhibit significant variation.

The highest performance was achieved by the complete PACCF framework with BK-Tree-based medical validation, reaching a CRA of 93.86% and a WRA of 78.52%. Compared with the confidence-only baseline, the full framework improved character-level accuracy by 20.28 percentage points and more than doubled word-level accuracy. The significant increase in WRA demonstrates that the combination of adaptive confidence fusion and efficient medical validation is particularly effective for correcting prescription-specific recognition errors. Overall, the ablation study confirms that PACCF benefits from the complementary interaction of confidence modeling, inter-model agreement analysis, statistical reliability estimation, medical lexical similarity, and BK-Tree-assisted validation.

Configuration	CRA (%)	WRA (%)
C (Confidence Only)	68.15	22.22
C + A (Agreement)	82.92	32.59
C + A + G (Gaussian)	85.12	26.67
C + A + G + M (Medical)	90.37	44.44
Full PACCF + BK-Tree	93.86	78.52

Table V. PACCF Ablation Analysis.

C. Contribution of Individual Factors

The contribution analysis presented in Table VI evaluates the importance of individual PACCF components using a leave-one-out experimental design. In each experiment, one reliability factor was removed while all remaining PACCF components were retained. The resulting performance degradation provides quantitative evidence regarding the contribution of each factor to the overall fusion framework.

The complete PACCF framework achieved the best overall performance, obtaining a Character Recognition Accuracy (CRA) of 93.86%, Word Recognition Accuracy (WRA) of 78.52%, Precision of 91.14%, Recall of 95.97%, and an F1 Score of 93.49%. These results serve as the baseline for assessing the contribution of individual factors.

When the Confidence Reliability (C) component was removed, CRA decreased from 93.86% to 93.35%, while WRA dropped from 78.52% to 77.78%. Precision, Recall, and F1 Score also declined slightly to 94.08%, 95.96%, and 94.56%, respectively, resulting in an overall accuracy drop of 0.51 percentage points. This indicates that confidence-based weighting contributes positively to the fusion process, although its impact is relatively modest when other reliability measures remain available.

Similarly, removing the Hybrid Agreement (A) component reduced CRA to 93.37% and WRA to 77.78%, with Precision, Recall, and F1 Score reaching 94.10%, 96.01%, and 94.59%, respectively. The corresponding accuracy drop of 0.49 percentage points suggests that inter-model consensus provides useful complementary information but is not the primary driver of recognition performance.

A substantially larger degradation was observed when the Gaussian Reliability (G) factor was excluded. In this setting, CRA decreased to 91.12% and WRA to 76.30%, while Precision, Recall, and F1 Score fell to 91.62%, 95.25%, and 92.63%, respectively. The resulting accuracy drop of 2.73 percentage points represents the largest decline among all ablation experiments, highlighting the critical role of statistical confidence modeling in identifying unreliable predictions and stabilizing the fusion process.

The removal of the Medical Similarity (M) factor also caused a significant reduction in performance. CRA declined to 91.56%, WRA decreased to 76.30%, and Precision, Recall, and F1 Score dropped to 92.25%, 94.41%, and 92.84%, respectively. The corresponding accuracy drop of 2.29 percentage points demonstrates the importance of domain-specific medical knowledge in distinguishing clinically plausible terms from visually similar but incorrect OCR outputs.

Overall, the leave-one-out analysis reveals that Gaussian Reliability and Medical Similarity are the most influential components of the PACCF framework, contributing the largest performance gains. In contrast, Confidence Reliability and Hybrid Agreement provide complementary improvements that further enhance robustness and consistency. These findings validate the design philosophy of PACCF, where statistical confidence estimation and domain-aware lexical validation work together to improve medical prescription recognition accuracy.

Configuration	CRA (%)	WRA (%)	Precision (%)	Recall (%)	F1 (%)	Accuracy Drop (%)
Full PACCF	93.86	78.52	91.14	95.97	93.49	0
Without Confidence Reliability (C)	93.35	77.78	94.08	95.96	94.56	0.51
Without Hybrid Agreement (A)	93.37	77.78	94.1	96.01	94.59	0.49
Without Gaussian Reliability (G)	91.12	76.3	91.62	95.25	92.63	2.73
Without Medical Similarity (M)	91.56	76.3	92.25	94.41	92.84	2.29

Table VI. Individual Factor Contribution Analysis

D. Retrieval Performance Analysis

The retrieval performance results presented in Table VII demonstrate the computational efficiency of the proposed BK-Tree-based medical validation framework. The objective of this experiment was to compare the query latency of BK-Tree retrieval against a conventional linear dictionary search strategy when matching OCR predictions to medicine names.

As shown in Table VII, the linear search approach required an average query time of 10.24 ms per lookup, whereas the BK-Tree retrieval mechanism achieved an average query time of only 0.26 ms. This corresponds to an approximate 33-fold reduction in retrieval latency. The substantial improvement is attributed to the hierarchical edit-distance indexing structure of the BK-Tree, which eliminates the need to compare a query against every medicine entry in the database.

From a computational perspective, the linear search strategy exhibits a retrieval complexity of (O(D)), where (D) denotes the size of the pharmaceutical vocabulary. In contrast, the BK-Tree performs edit-distance guided traversal and achieves an average retrieval complexity of $O(log\ log\ D\ )$.

Consequently, the performance advantage becomes increasingly significant as the medicine repository grows.

The experimental results indicate that BK-Tree indexing provides an efficient and scalable solution for domain-aware medical validation. By reducing retrieval overhead while preserving accurate approximate matching capabilities, the proposed approach enables real-time medicine validation and correction within the PACCF framework, making it suitable for large-scale healthcare document digitization systems.

Method	Avg. Time (ms)	Complexity
Linear Dictionary Search	10.2463813	O(D)
BK-Tree Retrieval	0.2606766	O(log D)

(Average-case retrieval complexity approaches $O(\log D)$.)

Table VII. Retrieval Performance Comparison.

Although the current experimental dataset is moderate in size, the observed retrieval acceleration demonstrates the suitability of the proposed framework for future deployment on substantially larger pharmaceutical vocabularies

E. Qualitative Examples

Table VIII. presents representative recognition examples obtained from challenging prescription images containing highly variable handwriting styles, character distortions, inconsistent spacing, and image quality degradations. These examples provide qualitative evidence of the effectiveness of the proposed PACCF framework beyond the quantitative metrics reported in previous sections.

As shown in Fig. 4, individual OCR engines frequently generate substantially different predictions for the same medicine name. For example, the medicine *Paracetamol* was recognized as *shreetavel*, *lasttlawnol*, and *paracetanol* by different OCR models, while *Amoxicillin* was transformed into fragmented or highly distorted outputs such as *amo vicilli n*, *amorialla*, and *amo xicallin*. Similar behavior was observed for *Aspirin* and *Losartan*, where several OCR engines produced outputs that exhibited little lexical similarity to the corresponding ground-truth medicine names.

Despite these inconsistencies, PACCF successfully recovered the correct medicine names in the majority of cases by combining confidence reliability, inter-model agreement, Gaussian consensus estimation, and medical lexical similarity within a unified fusion framework. The integration of medical-domain knowledge further enabled the correction of recognition errors that could not be resolved through confidence-based voting alone.

Notably, PACCF accurately recovered medicine names such as *Ibuprofen*, *Metformin*, *Omeprazole*, *Losartan*, and *Salbutamol* even when all individual OCR engines produced partially incorrect predictions. These examples demonstrate the ability of the proposed framework to exploit complementary recognition information and generate medically plausible outputs.

The most challenging example involved *Azithromycin*, where all OCR engines produced severely corrupted predictions with limited similarity to the ground truth. In this case, PACCF returned an uncertainty indication rather than forcing an incorrect correction. This behavior highlights the framework's ability to recognize ambiguous situations and avoid potentially unsafe medical substitutions. Overall, the qualitative results support the quantitative findings and demonstrate the robustness of PACCF for real-world prescription recognition scenarios.

Samples Images	**Ground_truth**	**Trocr**	**Parseq**	**Easyocr**	**Paddleocr**	**Paccf**
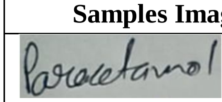	Paracetamol	paracetamol	shreetavel	lasttlawnol	paracetanol	paracetamol
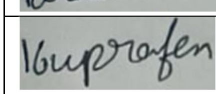	ibuprofen	-	louprofen	ibuprilen	-	ibuprofen
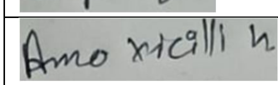	Amoxicillin	amo vicilli n	amorialla	-	amo xicallin	amoxicillin
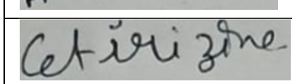	Cetirizine	cettigine	cetirizine	chisu ztne	cetirizine	cetirizine

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VI June 2026- Available at www.ijraset.com

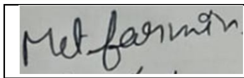	Metformin	met farments	mil farman	m janwz	melfarin	metformin
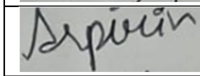	Aspirin	superior	aspucinin	jun	apoin	aspirin
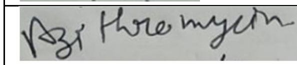	Azithrmycin	a z phone mycin	abikronych	kwshsum	azi thore myetn	[UNCERTAIN: abikronych]
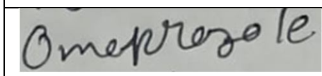	Omeprozole	o o mepreso e	omeproyele	le	omeprozele	omeprazole
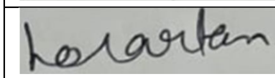	Losartan	resmortem	herantan	koaask	aartan	losartan
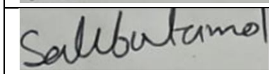	Salbutamol	salibutumol	salibiatemal	saxlubu lunol	salibutamol	salbutamol

(Examples are actual outputs produced during evaluation.)

Table VIII. Qualitative comparison of individual OCR engines and the proposed PACCF framework on handwritten prescription medicine recognition samples.

VII. DISCUSSION

A. Why PACCF Works

The experimental results demonstrate that the proposed Prescription-Aware Adaptive Confidence Fusion (PACCF) framework consistently outperforms individual OCR engines across all evaluation metrics. The primary reason for this improvement lies in the complementary recognition capabilities of the heterogeneous OCR ensemble. As observed in the qualitative examples, different OCR engines often produce distinct recognition errors when processing challenging prescription images. While one model may correctly recognize a medicine name, another may generate a partially corrupted prediction or an unrelated output. Consequently, selecting predictions solely on the basis of confidence scores may lead to unreliable decisions.

PACCF addresses this limitation by integrating multiple complementary reliability indicators within a unified adaptive fusion framework. Confidence Reliability captures model certainty, Hybrid Agreement measures consensus among OCR predictions, Gaussian Reliability evaluates statistical consistency with ensemble confidence behavior, and Medical Lexical Similarity introduces domain-specific pharmaceutical knowledge into the decision-making process. The adaptive Softmax weighting mechanism dynamically determines the contribution of each factor, enabling data-driven fusion without manually tuned weighting coefficients.

The ablation and factor-contribution analyses provide further evidence supporting the effectiveness of the proposed design. In particular, Gaussian Reliability and Medical Similarity produced the largest performance contributions, indicating that statistical confidence modeling and domain-aware validation play a critical role in improving recognition robustness. The substantial increase in Word Recognition Accuracy from 31.11% for the strongest baseline model to 78.52% for PACCF further demonstrates that combining ensemble diversity with medical-domain validation is highly effective for correcting prescription-specific recognition errors.

B. Failure Cases and Limitations

Despite its strong performance, PACCF remains subject to several limitations. The framework assumes that at least one OCR engine generates a reasonably informative candidate prediction. When all OCR engines fail simultaneously due to extremely poor handwriting, severe image blur, low illumination, occlusions, or significant image degradation, the fusion module has insufficient information to reliably infer the correct medicine name.

A representative example is the recognition of *Azithromycin*, where all OCR engines produced highly corrupted outputs with limited lexical similarity to the ground truth. In such situations, PACCF intentionally returns an uncertainty indication rather than forcing an incorrect correction. Although this conservative behavior reduces the risk of unsafe medical substitutions, it also highlights the dependence of the framework on the quality of candidate predictions generated by the OCR ensemble.

Another limitation arises from the medicine knowledge repository. If a medicine name is absent from the database, appears as a rare regional brand name, or is represented using uncommon abbreviations, the lexical validation stage may be unable to provide an appropriate correction. Furthermore, lexical similarity measures alone may not fully distinguish between pharmaceutical terms that are visually similar yet clinically different.

From a computational perspective, although BK-Tree indexing substantially reduces retrieval overhead, the overall execution time remains dominated by the inference cost of multiple OCR engines. Consequently, processing latency may increase when large volumes of prescription images are analyzed simultaneously or when computational resources are limited.

Another limitation of the present study is the relatively small size of the experimental dataset, which consists of 137 handwritten prescription images. Although the dataset was collected under diverse real-world conditions and contains substantial handwriting variability, it may not fully represent the wide range of prescription styles encountered in large-scale healthcare environments. Therefore, additional evaluation on larger and more diverse prescription datasets is required to further assess the generalizability and robustness of the proposed PACCF framework.

C. Practical Deployment Considerations

The modular architecture of PACCF makes it well suited for practical healthcare document digitization systems. Since the fusion mechanism operates independently of any specific OCR engine, individual recognition models can be upgraded, replaced, or extended without modifying the core PACCF framework. This design improves maintainability and facilitates compatibility with future OCR technologies.

The BK-Tree-based validation module enables efficient approximate matching against large pharmaceutical vocabularies, supporting scalable deployment in real-world healthcare environments. In addition, the adaptive fusion strategy eliminates the need for manually tuned weighting parameters, improving portability across datasets exhibiting different handwriting styles, imaging conditions, and prescription formats.

Potential application areas include pharmacy automation systems, prescription digitization platforms, electronic health record generation, healthcare document management systems, and clinical information extraction pipelines. Future research will focus on expanding the pharmaceutical knowledge repository, incorporating contextual prescription understanding, evaluating larger and more diverse datasets, and investigating learning-based fusion mechanisms capable of automatically modeling reliability relationships from large-scale prescription corpora.

VIII. CONCLUSION AND FUTURE WORK

This paper presented the Prescription-Aware Adaptive Confidence Fusion (PACCF) framework for robust medical prescription text recognition under challenging real-world conditions. The proposed framework addresses the limitations of conventional single-engine OCR systems by integrating a heterogeneous OCR ensemble consisting of TrOCR, PARSeq, EasyOCR, and PaddleOCR within a unified adaptive fusion architecture. Unlike traditional confidence-based voting methods, PACCF incorporates four complementary reliability indicators—Confidence Reliability, Hybrid Agreement, Gaussian Consensus Reliability, and Medical Lexical Similarity—to evaluate candidate predictions and perform adaptive decision fusion.

To further enhance domain-specific recognition accuracy, a BK-Tree-based medical validation module was introduced for efficient approximate retrieval and correction of pharmaceutical terminology. The proposed validation strategy enables scalable medicine matching while reducing retrieval complexity from (O(D)) to (O(\log D)), thereby improving computational efficiency for large pharmaceutical vocabularies.

Experimental evaluation conducted on a challenging dataset of 137 handwritten prescription images demonstrated the effectiveness of the proposed framework. PACCF achieved a Character Recognition Accuracy of 93.86% and a Word Recognition Accuracy of 78.52%, substantially outperforming all individual OCR engines. Ablation studies further confirmed the contribution of adaptive fusion and domain-aware validation components, while qualitative analysis demonstrated the framework's ability to recover medically plausible predictions from highly inconsistent OCR outputs.

Overall, the results indicate that combining ensemble diversity, adaptive confidence modeling, statistical reliability estimation, and medical-domain knowledge provides a practical and effective solution for prescription digitization. The proposed framework offers a promising approach for healthcare document processing applications where recognition accuracy and reliability are critical.

Future work will focus on expanding the pharmaceutical knowledge repository, incorporating contextual prescription understanding, and investigating learning-based fusion mechanisms capable of automatically optimizing reliability relationships from data. Although the proposed framework demonstrated strong performance on a dataset of 137 handwritten prescription images, further evaluation on larger and more diverse prescription datasets is required to comprehensively assess its generalization capability across different handwriting styles, imaging conditions, and clinical settings. Additional research may explore lightweight deployment architectures, multilingual prescription recognition, and multimodal healthcare document understanding systems to support large-scale clinical digitization environments.

IX. ACKNOWLEDGMENT

The author conducted this research independently and would like to acknowledge the contributions of the open-source research community, particularly the developers of TrOCR, PARSeq, EasyOCR, and PaddleOCR, whose publicly available tools and resources enabled this study. No external funding or institutional support was received for this work.

REFERENCES

Core OCR and Transformer References

[1] M. Li et al., "TrOCR: Transformer-Based Optical Character Recognition with Pre-Trained Models," Proc. AAAI Conf. Artif. Intell., vol. 37, no. 11, pp. 13094–13102, 2023.
[2] D. Bautista and R. Atienza, "Scene Text Recognition with Permuted Autoregressive Sequence Models," in Proc. Eur. Conf. Comput. Vis. (ECCV), 2022.
[3] A. Vaswani et al., "Attention Is All You Need," in Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 30, 2017.
[4] A. Dosovitskiy et al., "An Image Is Worth 16×16 Words: Transformers for Image Recognition at Scale," in Proc. Int. Conf. Learn. Represent. (ICLR), 2021.
[5] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks," in Proc. ICML, pp. 369–376, 2006.
[6] R. Smith, "An Overview of the Tesseract OCR Engine," in Proc. ICDAR, pp. 629–633, 2007.
[7] T. Coquenet, C. Chatelain, and T. Paquet, "End-to-End Handwritten Paragraph Text Recognition Using a Vertical Attention Network," IEEE Trans. Pattern Anal. Mach. Intell., vol. 45, no. 1, pp. 508–524, 2023.

Medical Prescription OCR

[8] U. Ali et al., "Leveraging Deep Learning with Multi-Head Attention for Accurate Extraction of Medicine from Handwritten Prescriptions," arXiv preprint arXiv:2412.18199, 2024.
[9] H. Ullah et al., "Enhancing Handwritten Prescription Recognition with AI-Driven OCR Frameworks," J. Comput. Biomed. Informat., 2025.
[10] T. Mankash et al., "MIRAGE: Multimodal Identification and Recognition of Annotations in Indian General Prescriptions," arXiv preprint arXiv:2410.09729, 2024.
[11] S. R. Gudi, "Enhancing OCR Accuracy in Healthcare Prescription Processing Using Artificial Neural Networks," Eur. J. Artif. Intell. Mach. Learn., vol. 4, no. 6, 2025.

OCR Ensemble and Fusion

[12] A. Goel and A. Kaushik, "Prescription-to-Text Conversion Using Ensemble Faster R-CNN and OCR," 2023.
[13] E. Hassan et al., "Medical Prescription Recognition Using Machine Learning," 2021.
[14] J. Kittler, M. Hatef, R. P. W. Duin, and J. Matas, "On Combining Classifiers," IEEE Trans. Pattern Anal. Mach. Intell., vol. 20, no. 3, pp. 226–239, 1998.