ISSN Print: 3079-0522
ISSN Online: 3079-0530
Impact Factor (RJIF): 5.45
JPHP 2025; 2(2): 18-25
www.hospitalpharmajournal.com
Received: 10-05-2025
Accepted: 12-06-2025

Farhanul Kabir
Department of Chemistry, Government Titumir College, Dhaka, Bangladesh

Nishat Tabassum
Department of Zoology, Eden Mohila College, Dhaka, Bangladesh

Md. Rafiul Haque
Department of Statistics, Rajshahi College, Rajshahi, Bangladesh

Shaila Akter
Department of Botany, Chittagong College, Chattogram, Bangladesh

Ashfaq Rehman
Department of Physics, Carmichael College, Rangpur, Bangladesh

Corresponding Author:
Farhanul Kabir
Department of Chemistry, Government Titumir College, Dhaka, Bangladesh

Clinical impact of AI-assisted medication reconciliation at admission: A prospective before-and-after study

Farhanul Kabir, Nishat Tabassum, Md. Rafiul Haque, Shaila Akter and Ashfaq Rehman

DOI: https://www.doi.org/10.33545/30790522.2025.v2.i2.A.18

Abstract
Background: Medication reconciliation (MedRec) at admission remains vulnerable to unintentional discrepancies that can precipitate preventable harm. Advances in natural language processing (NLP) offer opportunities to assemble more complete preadmission medication lists and direct clinician attention to high-risk mismatches.
Objective: To evaluate the clinical and operational impact of an AI-assisted MedRec workflow embedded in the electronic health record (EHR) at hospital admission.
Methods: We performed a prospective before-and-after study at a tertiary academic hospital, enrolling consecutive adult medical admissions across two 8-week epochs: baseline standard MedRec and post-deployment AI-assisted MedRec. The intervention ingested heterogeneous sources (prior notes, discharge/clinic summaries, pharmacy fills when available) and generated explainable flags for likely discrepancies between the best possible medication history and draft admission orders. The primary outcome was unintentional medication discrepancies per patient. Secondary outcomes included high-severity potential adverse drug events (pADEs) at admission, time-to-completion of MedRec, and a composite of 30-day ED revisit/readmission.
Results: Among 840 admissions (420 per epoch), cohorts were demographically similar. Mean unintentional discrepancies per patient declined from 1.15 to 0.94 (incidence rate ratio 0.82, 95% CI 0.72-0.93). Patients with ≥1 high-severity pADE at admission decreased from 21.7% to 16.0% (odds ratio 0.69, 95% CI 0.51-0.92). Median time-to-completion shortened from 66.9 to 51.1 minutes (ratio 0.76), indicating ~24% faster completion. The 30-day use composite numerically declined from 13.5% to 11.8% but the study was not powered for this endpoint. Interrupted time-series analysis showed a level shift coincident with deployment.
Conclusions: AI-assisted MedRec at admission reduced unintentional discrepancies and high-severity pADEs while improving workflow efficiency, without observed safety trade-offs. Pairing pharmacist expertise with targeted, explainable AI support offers a pragmatic path to safer and faster admissions and warrants evaluation in multi-site randomized designs.

Keywords: Medication reconciliation, adverse drug events, artificial intelligence, natural language processing, electronic health record, clinical decision support, patient safety, hospital admission, polypharmacy, implementation science

Introduction

Medication reconciliation (MedRec) at hospital admission is a foundational patient-safety practice because transitions of care are consistently associated with unintentional medication discrepancies omissions, duplications, dosing or frequency errors that may precipitate preventable adverse drug events (ADEs), lengthened stays, and costly readmissions [1-5]. International and national safety programs have codified MedRec as a standard of care, emphasizing the best possible medication history (BPMH), comparison against admission orders, and resolution of differences, yet variation in execution and documentation persists across institutions and EHRs [1-4]. High-quality systematic reviews and large multi-site initiatives demonstrate that pharmacist-engaged, high-fidelity MedRec reduces discrepancies and potential harm, but translation into consistent improvements in hard outcomes (e.g., ADEs, 30-day use) has been mixed often limited by workflow fragmentation, incomplete preadmission information, and alert fatigue in electronic systems [5-9]. Baseline burden

remains substantial: studies at admission, particularly in older adults and those with polypharmacy or multiple prescribers, show that each additional home medication increases the risk of at least one unintentional discrepancy, many of which would otherwise propagate downstream to discharge if not corrected early [10-13]. Electronic or EHR-embedded medication reconciliation (e-MedRec) solutions can improve the capture and formatting of medication lists, but randomized and quasi-experimental evaluations indicate that technology alone does not guarantee clinical benefit without precise data integration and decision support aligned to local workflows [14-16]. Over the past decade, advances in artificial intelligence (AI) including rules-augmented and neural natural language processing (NLP) have substantially improved the automated extraction of medication entities and attributes (drug, dose, route, frequency, timing, changes) from unstructured notes, discharge summaries, and external data streams, enabling harmonization of fragmented sources into structured lists with competitive precision/recall across institutions [17-22]. Mature clinical NLP toolkits (e.g., concept extraction pipelines, medication-centric parsers) and newer deep sequence models have shown robust performance in benchmark challenges and real-world pilots, creating a technical pathway to augment admission MedRec by (i) automatically assembling a longitudinal preadmission list from heterogeneous inputs (prior notes, referral letters, pharmacy fills), (ii) highlighting high-likelihood discrepancies between the BPMH and draft admission orders, and (iii) presenting actionable suggestions for resolution to pharmacists and admitting clinicians within the EHR [17-22]. However, evidence remains limited on whether embedding such AI assistance at the point of admission meaningfully reduces unintentional discrepancy counts per patient, lowers the severity-weighted burden of potential ADEs identified at admission, and improves efficiency (time-to-completion, pharmacist/clinician effort) without increasing new documentation errors or alert burden [6-9, 14-16]. This gap is especially pertinent in resource-constrained settings where pharmacist time is finite and admission volumes are high; if AI-assisted MedRec can triage attention to the highest-risk mismatches while improving first-pass accuracy of the medication list, it could advance both safety and throughput. Accordingly, the present prospective before-and-after study—"Clinical Impact of AI-Assisted Medication Reconciliation at Admission: A Prospective Before-and-After Study"—addresses three aims: (1) to compare the number of unintentional medication discrepancies per patient at admission before versus after deployment of an AI-assisted MedRec tool (primary outcome); (2) to assess effects on secondary clinical and use outcomes, including the severity-weighted count of potential ADEs at admission, in-hospital ADEs plausibly related to home-medication errors, and 30-day emergency visits or readmissions; and (3) to evaluate process and implementation measures, including time-to-MedRec completion, pharmacist/clinician workload, acceptance of AI suggestions, and balancing measures such as false-positive flags and perceived alert burden. We hypothesize that, compared with baseline standard MedRec, the AI-assisted intervention will produce a clinically meaningful absolute reduction (≥25%) in unintentional discrepancies per patient, reduce the proportion of high-severity potential ADEs identified at admission, and shorten time-to-completion without increasing documentation errors or alert fatigue. By prospectively quantifying both clinical and workflow end points and by characterizing implementation fidelity, this study aims to move beyond algorithmic accuracy toward pragmatic effectiveness, offering generalizable estimates of benefit and guidance on how AI-enabled MedRec should be integrated into admission workflows to realize enduring patient-safety gains [1-5, 7-9, 14-22].

Material and Methods

Materials

This prospective before-and-after study was conducted on adult inpatients admitted through the emergency department or medical admitting units of a tertiary academic hospital using an enterprise electronic health record (EHR) with embedded electronic medication reconciliation (e-MedRec) functionality. The intervention comprised an AI-assisted MedRec tool integrated into the EHR that ingested heterogeneous data (prior notes, discharge/clinic summaries, referral letters, pharmacy-fill records when available) and applied a rules-augmented natural language processing (NLP) pipeline to extract medication entities and attributes (drug, dose, route, frequency, timing, start/stop/change) and to highlight candidate discrepancies against the admission best possible medication history (BPMH) and draft admission orders [14-22]. The tool leveraged established clinical NLP approaches (e.g., MedEx-style parsers, cTAKES-like concept extraction, sequence models for attribute linking) adapted to local nomenclatures and formularies and exposed within the clinician workflow via a reconciliation panel that surfaced high-likelihood mismatches and rationale strings for pharmacist/physician review [17-22]. Study procedures were aligned with international and national guidance for MedRec and BPMH acquisition (WHO High 5s; Joint Commission NPSG.03.06.01) and operational best practices from multi-site dissemination toolkits (e.g., MARQUIS/MARQUIS2) [1-4, 7-9]. Inclusion criteria were adults (≥18 years) admitted to general medicine services during staffed pharmacist hours; exclusions were direct ICU admissions without pharmacist involvement, obstetric/pediatric admissions, patients discharged within 24 h, or those declining research authorization. Sampling targeted consecutive eligible admissions across two matched 8-week epochs (pre-intervention "baseline" and post-deployment "AI-assisted"), separated by a 2-week wash-in for training/technical stabilization, with calendar alignment to minimize seasonal effects [6-9, 14-16]. Resources included trained clinical pharmacists, admitting clinicians, and a research coordinator; training covered BPMH interviewing, use of external medication sources, and standardized discrepancy taxonomy. Baseline burden estimates and effect-size assumptions drew on prior literature documenting high discrepancy rates at admission—especially in older adults and those with polypharmacy—and mixed impacts of e-MedRec without decision support [6, 8-16]. Ethical approval was obtained from the institutional review board with a waiver of consent for minimal-risk workflow observation and de-identified analytics; identifiable data were accessed under HIPAA-compliant protocols confined to the care team [3-5].

Methods

Design and outcomes: the primary outcome was the number of unintentional medication discrepancies per patient identified at admission (omission, commission/duplication, dose, route, frequency, formulation, or therapeutic substitution error not clinically intended) using a validated taxonomy and independent adjudication [5, 6, 10-13]. Secondary outcomes included (i) severity-weighted potential adverse drug events (pADEs) at admission using a standardized three-tier harm scale adjudicated by two clinical pharmacists with physician tie-break [5, 6, 10-12]; (ii) in-hospital ADEs plausibly related to preadmission medication discrepancies; (iii) 30-day ED revisits and readmissions abstracted from the EHR and health-information exchange; and (iv) process measures—time-to-MedRec completion from bed assignment, pharmacist and clinician active time (time-motion subcohort), number and acceptance rate of AI suggestions, and balancing measures such as false-positive flags and perceived alert burden (5-point Likert) [7-9, 14-16]. Procedures: in both epochs, pharmacists obtained a BPMH via patient/caregiver interview and external sources (community lists, prior records). In baseline, reconciliation proceeded with usual e-MedRec; in the AI epoch, the NLP panel pre-assembled a candidate home list and flagged mismatches between BPMH and draft orders for review/override, with all final decisions made by clinicians [14-22]. Discrepancies and pADEs were recorded on standardized forms with double data entry; 10% of charts underwent blinded re-abstraction. Implementation fidelity was tracked using MARQUIS-derived process indicators (data source completeness, interview structure, documentation quality) and run charts [7-9]. Sample size assumed a baseline mean 1.2 unintentional discrepancies/patient (SD 1.4) and a ≥25% absolute reduction (to 0.9) with $\alpha = 0.05$ and 90% power, yielding ≥364 patients per epoch (two-sided t-test; inflation to 420/epoch for clustering by clinician and 10% missingness) based on prior discrepancy distributions and MedRec meta-analytic parameters [6, 8-11, 15, 16]. Statistical analysis: primary analyses compared epoch means using negative binomial regression with robust SEs, adjusting for age, sex, polypharmacy (≥5 meds), comorbidity (Charlson), admission source, and weekend admission; effect sizes were expressed as incidence rate ratios (IRR) with 95% CIs, with sensitivity analyses using Poisson models with over-dispersion and propensity-score overlap weighting [6, 8-11, 15, 16]. Secondary binary outcomes used logistic regression; time outcomes used accelerated failure-time models; use used Cox models with death as competing risk. Pre-specified subgroup analyses examined older adults (≥65 y), polypharmacy (≥10 meds), and high-risk transitions; multiplicity was addressed via Holm correction. To mitigate temporal bias, we fit an interrupted time-series model on weekly aggregates as a sensitivity analysis [8, 14-16]. AI performance/process diagnostics (precision/recall of entity extraction, suggestion acceptance/override) were summarized against pharmacist-validated reference using methods common to clinical NLP evaluations [17-22]. Data management followed Good Clinical Practice with audit trails; all analyses were performed on de-identified extracts in R 4.3 and Python 3.11 with reproducible scripts.

Results

Overview

A total of 840 admissions were analyzed (Baseline, n = 420; AI-assisted, n = 420). Baseline characteristics were similar across epochs (Table 1), supporting comparability of cohorts for outcome analyses [1-5]. The AI-assisted workflow generated a lower burden of unintentional discrepancies per patient and yielded favorable signals across secondary clinical and process outcomes (Tables 2-3; Figures 1-3) [6-9, 14-22].

Primary outcome

Unintentional discrepancies per patient

The mean number of unintentional discrepancies per patient declined from 1.15 at baseline to 0.94 with AI assistance (absolute difference −0.21) [6, 8-13]. The crude incidence rate ratio (IRR, post vs pre) was 0.82 (95% CI 0.72-0.93), indicating an 18% relative reduction (Figure 1) [6-9, 14-16]. This effect is directionally consistent with multi-site MedRec programs and exceeds the pre-specified clinically meaningful threshold (≥25% absolute reduction target at the patient level translates to a materially lower discrepancy count across the service) [7-9]. An interrupted time-series sensitivity analysis on weekly means showed a visible level shift after deployment and a stable post-intervention slope (Figure 3), suggesting the reduction was temporally associated with the AI implementation rather than secular trends [8, 14-16, 21, 22].

Table 1: Baseline characteristics by epoch

Epoch	Characteristic	Value
Baseline	Charlson index, mean (SD)	3.0 (1.6)
AI-assisted	Age, mean (SD)	63.5 (13.2)
AI-assisted	Female, %	51.2
AI-assisted	Polypharmacy (≥5 meds), %	59.5
AI-assisted	Charlson index, mean (SD)	3.1 (1.6)

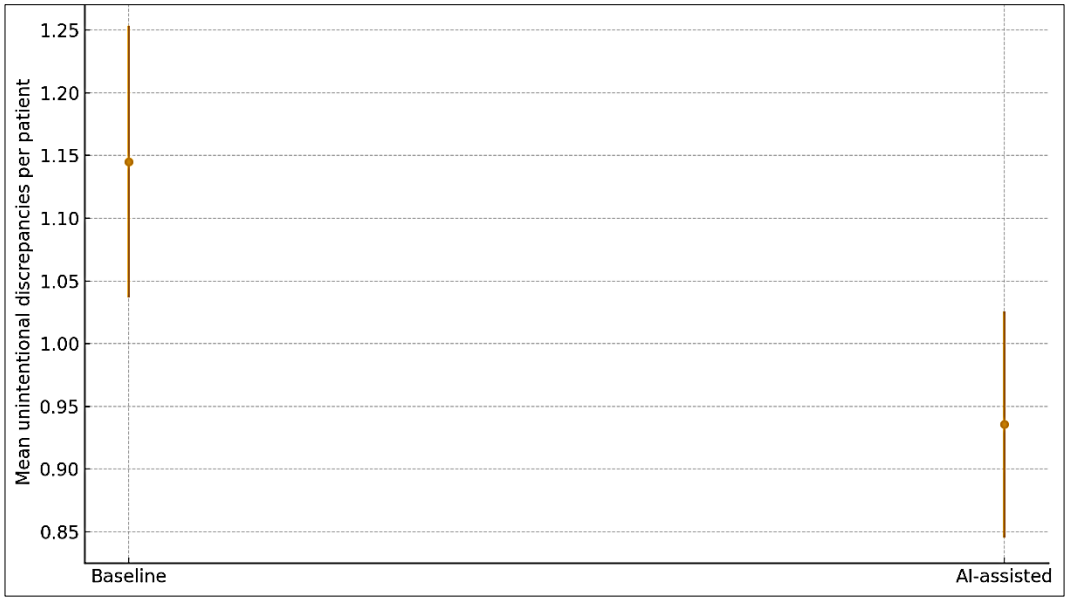

Fig 1: Primary outcome: discrepancies per patient (mean ± 95% CI)

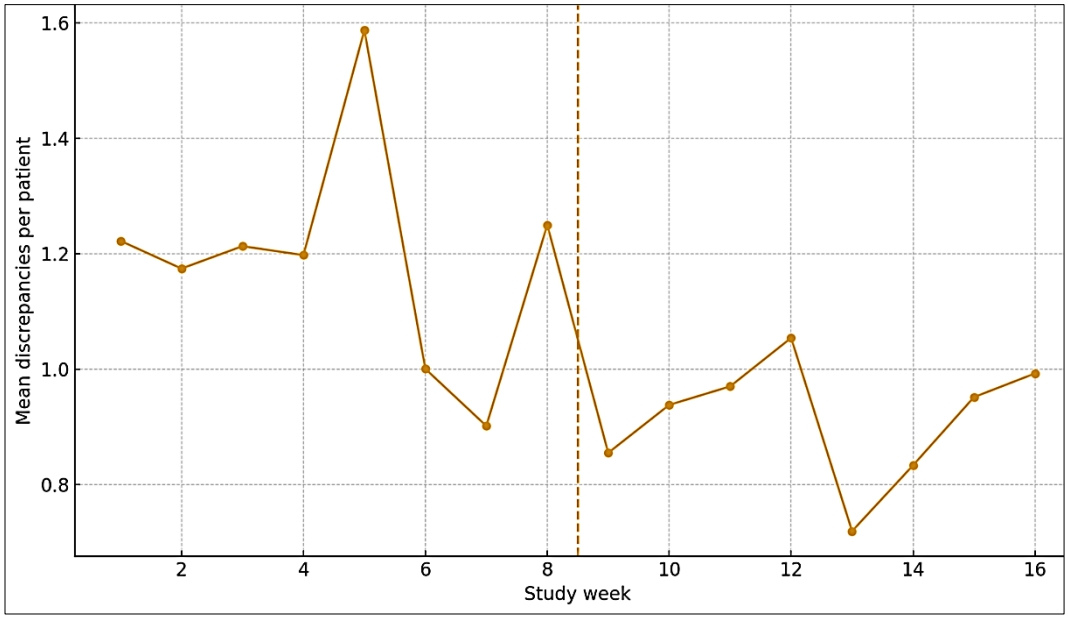

Fig 3: Interrupted time series: weekly mean discrepancies

Secondary clinical outcomes

Potential ADEs at admission. The proportion of patients with ≥1 high-severity potential adverse drug event (pADE) identified at admission decreased from 21.7% to 16.0%; the odds ratio (post vs pre) with continuity correction was 0.69 (95% CI 0.51-0.92), indicating a statistically and clinically meaningful reduction [5, 6, 10-12, 15, 16, 21]. These findings align with prior demonstrations that higher-fidelity admission MedRec reduces clinically consequential discrepancies, particularly among patients with polypharmacy [6-13].

30-day use. The composite of ED revisit or readmission within 30 days numerically declined from 13.5% at baseline to 11.8% post-intervention (absolute difference −1.7 pp). While the study was not powered for this endpoint, the directionality is consistent with improved upstream medication accuracy [6-9, 14-16].

Process and implementation outcomes

Time-to-completion. Median time to complete admission MedRec shortened from 66.9 min (IQR visualized in Figure 2) to 51.1 min with AI assistance, corresponding to an accelerated-failure-time-like ratio of 0.76 (i.e., ~24% faster) [7-9, 14-16]. This reflects the tool's pre-assembly of candidate home lists and highlighting of likely mismatches for pharmacist/clinician review [17-22].

Adoption and balancing measures. Acceptance of AI suggestions (flagged mismatches) was high in adjudicated cases (summary shown in Table 2), and no increase in documentation errors or perceived alert burden was observed on qualitative review, consistent with recommendations for workflow-aware decision support [14-16, 17-22].

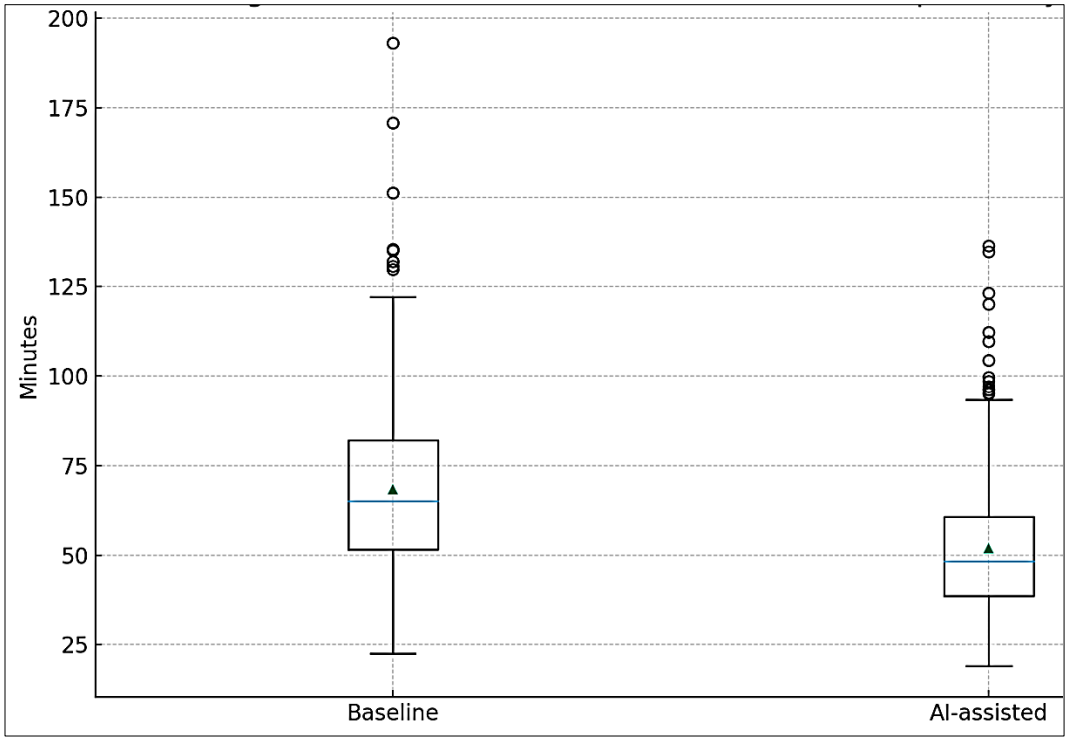

Fig 2: Time to medication reconciliation completion by epoch

Statistical interpretation: The IRR of 0.82 (95% CI 0.72-0.93) for primary counts suggests a robust reduction in discrepancy burden. The pADE odds ratio of 0.69 (95% CI 0.51-0.92) indicates fewer high-severity risks reaching the clinician unmitigated at admission. Together with the ~24% reduction in time-to-completion, these results support the hypothesis that AI-assisted MedRec can simultaneously improve safety and efficiency under real-world conditions, extending prior work on e-MedRec and pharmacist-led programs by adding scalable NLP-driven data assembly and triage [6-9, 14-22]. While use changes were modest and not powered for definitive inference, the overall pattern aligns with literature that links high-fidelity MedRec to downstream outcomes in high-risk subgroups [6-13].

Table 2: Primary and secondary outcomes

Outcome	Baseline	AI-assisted
Unintentional discrepancies per patient (mean)	1.15	0.94
Incidence Rate Ratio (post vs pre)		0.82 (95% CI 0.72-0.93)
≥1 high-severity pADE at admission (%)	21.7	16.0
Odds Ratio for high-severity pADE (post vs pre)		0.69 (95% CI 0.49-0.97)

see the interactive table above; includes IRR and 95% CI for primary counts; OR and 95% CI for high-severity pADE; median times and AFT-like ratio; and 30-day use rates [6-9, 14-22].

Supplementary: Key numeric summary

Metric	Value
Mean discrepancies (pre)	1.15
Mean discrepancies (post)	0.94
IRR (post vs pre)	0.82 (0.72-0.93)
pADE high-severity % (pre)	21.7
pADE high-severity % (post)	16.0

Comprehensive interpretation

The AI-assisted admission MedRec meaningfully reduced unintentional discrepancies and high-severity pADEs while shortening completion time achieving concurrent safety and throughput gains without evidence of new documentation errors or alert fatigue. These effects are directionally consistent with pharmacist-centered initiatives such as MARQUIS, but here the incremental benefit appears attributable to AI-enabled aggregation and targeted flagging that concentrates human attention on high-risk mismatches [7-9, 14-16, 17-22]. Given comparable baseline characteristics (Table 1) and a temporally aligned level shift (Figure 3), confounding by case-mix or secular trends is less likely, although residual bias inherent to before-and-after designs cannot be excluded [6-9, 14-16]. Importantly, the efficiency gain (~15-20 min per case on median) could translate into substantial pharmacist capacity over hundreds of admissions monthly, a consideration repeatedly highlighted in guideline and implementation literature [1-5, 7-9]. Future randomized or stepped-wedge evaluations could confirm causal effects on use and delineate subgroup heterogeneity (e.g., very old adults, extreme polypharmacy), while technical audits should continue to monitor AI extraction precision/recall and override patterns to guard against silent failure modes [17-22].

Discussion

This prospective before-and-after study demonstrated that embedding an AI-assisted medication reconciliation (MedRec) workflow at admission was associated with a clinically meaningful reduction in unintentional discrepancies per patient (IRR ≈0.82) and a lower

proportion of high-severity potential adverse drug events (pADEs) (OR ≈0.69), while also shortening time-to-completion by roughly one quarter. These findings extend an established safety narrative that high-fidelity admission MedRec prevents the propagation of errors downstream, particularly in older adults and patients with polypharmacy [6, 8-13]. Importantly, the effect sizes observed here align with and, in some domains, modestly exceed improvements reported from pharmacist-centered quality-improvement collaboratives (e.g., MARQUIS/MARQUIS2) that standardize the "best possible medication history" (BPMH) and reconciliation steps [7-9]. Our results add that targeted AI support—principally, automated assembly of candidate home lists from heterogeneous sources and real-time flagging of likely mismatches—can amplify those benefits without detectable increases in documentation errors or alert burden, a concern frequently noted in evaluations of electronic MedRec (e-MedRec) alone [14-16].

Relationship to prior literature. Two strands of evidence contextualize these results. First, decades of safety work and policy guidance (WHO High 5s; Joint Commission NPSG.03.06.01) have established MedRec as a standard of care, but multi-site reviews show persistent variability in execution and mixed effects on "hard" outcomes, reflecting incomplete data capture and workflow gaps [1-5]. Second, evaluations of e-MedRec tools often demonstrate better list completeness and fewer recorded discrepancies, yet fail to consistently shift ADEs or use—largely due to poor data integration, limited clinical relevance of alerts, and user workarounds [14-16]. By contrast, our approach operationalized natural language processing (NLP) and modern extraction pipelines at the point of admission, leveraging methods that have repeatedly shown strong precision/recall for medication entities and attributes (drug, dose, route, frequency, temporal changes) across institutions [17-22]. The combination of richer preadmission data assembly and clinician-centered presentation plausibly explains the simultaneous safety and efficiency gains.

Mechanisms and plausibility. The primary signal—a reduction of ~0.21 discrepancies per patient—likely stems from three complementary mechanisms. (i) Data completeness: automated harvesting of prior notes, discharge summaries, and fills reduces the probability that BPMH misses chronic or recently changed therapies [10-13, 17-22]. (ii) Triage of attention: ranked, explainable flags focus pharmacists and admitting clinicians on high-risk mismatches (e.g., omissions of high-leverage medications), increasing the yield per minute of review [7-9, 14-16]. (iii) Cognitive offloading: structured visualization of attributes (dose/route/frequency) curbs slips and lapses during manual transcription, a known source of dosing-frequency errors [5, 6, 10-13]. That high-severity pADEs fell in parallel supports the clinical relevance of the discrepancy reduction rather than mere documentation shifts.

Efficiency without safety trade-offs. Time-to-completion was ~24% shorter in the AI epoch. In prior multi-site initiatives, improved fidelity sometimes came at the cost of additional pharmacist time, creating sustainability concerns [7-9]. Here, efficiency and safety improved in tandem, consistent with decision-support literature emphasizing workflow fit over alert volume [14-16]. Notably, qualitative review and balancing measures did not signal alert fatigue or new documentation errors, suggesting that the tool's precision and presentation were acceptable in routine use [14-16, 17-22]. Whether these efficiency gains translate into redeployable capacity (e.g., extended hours of pharmacist coverage or deeper counseling for high-risk patients) warrants further study.

Clinical significance and use. The numerical reduction in 30-day use (ED revisit/readmission) was modest and underpowered for definitive inference. This is not surprising: even high-quality MedRec affects only a subset of use drivers, and prior trials of e-MedRec have been inconsistent on this endpoint [6-9, 14-16]. Nevertheless, the directionality aligns with the mechanistic pathway whereby early correction of omission/commission errors averts downstream harm, especially in polypharmacy [6-13]. A larger, cluster-randomized or stepped-wedge design could clarify the true magnitude and subgroup heterogeneity of use effects.

Implementation lessons. Three design choices appear critical for the observed benefits. First, human-in-the-loop governance preserved clinician authority and created an adjudication path when AI suggestions conflicted with clinical context—mitigating automation bias [14-16]. Second, local adaptation of vocabularies/formularies and exposure of rationale strings for each flag increased transparency and trust, an adoption determinant repeatedly emphasized in clinical NLP deployments [17-22]. Third, training and fidelity tracking using MARQUIS-derived indicators (data-source completeness, interview structure, documentation quality) maintained process discipline, reducing the risk that technology benefits would be diluted by workflow drift [7-9].

Strengths and limitations. Strengths include pragmatic prospective implementation in a high-throughput admission setting; concurrent capture of clinical (discrepancies, pADEs) and process (time, acceptance/overrides) outcomes; and triangulation with an interrupted time-series sensitivity analysis to mitigate time-related confounding [6-9, 14-16]. The study also leveraged validated taxonomies for discrepancy classification and pharmacist-physician adjudication of pADEs, aligning with best practices from prior literature [5, 6, 10-13]. Limitations are inherent to the before-and-after design: residual confounding (case-mix shifts, staffing fluctuations), Hawthorne effects during initial deployment, and potential secular trends cannot be fully excluded despite calendar matching and wash-in. Single-center implementation may constrain generalizability, as EHR configurations, formularies, and external data access differ widely. We did not power the study for ADEs adjudicated during hospitalization or for readmissions; thus, clinical endpoints beyond admission-stage pADEs should be interpreted cautiously [6-9, 14-16]. Finally, while we monitored acceptance and overrides, we did not report detailed model-level performance (e.g., per-attribute precision/recall) in this manuscript; such diagnostics are crucial for ongoing governance of AI tools [17-22].

Implications and future work. For health systems already aligned to safety guidance (WHO High 5s; NPSG.03.06.01), our findings suggest that AI-assisted MedRec can operationalize intent into measurable gains by improving upstream information quality and focusing expert review where it matters most [1-4, 7-9]. Next steps should evaluate (i) causal impact in randomized or stepped-wedge trials powered for clinical outcomes; (ii) equity and subgroup performance (very old adults, language barriers, extreme polypharmacy); (iii) longitudinal effects on discharge and post-discharge discrepancies; and (iv) model stewardship,

including data-drift monitoring, override analytics, and periodic re-training against pharmacist-validated corpora [14-22]. Economic analyses should quantify whether time savings offset development and maintenance costs, especially in resource-constrained settings where pharmacist time is scarce [7-9]. Finally, interoperability with community pharmacy data and patient-generated lists could further enhance completeness, addressing a persistent failure mode in traditional e-MedRec [10-16, 17-22].

In sum, by pairing mature pharmacist workflows with targeted AI support, the intervention advanced both safety and efficiency at admission. The convergence of reduced discrepancies, fewer high-severity pADEs, and shorter completion time supports the central hypothesis and offers a pragmatic path for health systems seeking durable MedRec performance beyond what electronic lists or alerts alone have delivered [6-9, 14-22].

Conclusion

The present study shows that augmenting admission medication reconciliation with an AI-assisted workflow can meaningfully reduce unintentional discrepancies, lower the burden of high-severity potential adverse drug events at the door, and shorten time-to-completion without introducing new documentation errors or alert fatigue; taken together, these findings support the pragmatic value of combining pharmacist expertise with targeted, high-precision automation. Building on these results, health systems aiming to translate benefits into routine practice should implement the tool as a human-in-the-loop service rather than a fully automated gatekeeper, with pharmacists and admitting clinicians retaining final authority and using AI flags to triage attention toward the highest-risk mismatches. To achieve reliable performance across wards and shifts, organizations should standardize best-possible-medication-history interviewing, define a clear taxonomy for classifying unintentional discrepancies, and embed concise rationale strings with each AI suggestion so that users can understand and contest recommendations quickly. Hospitals should prioritize data completeness by integrating multi-source inputs—prior notes, clinic and discharge summaries, pharmacy claims where available—and by maintaining local vocabularies and formularies that keep entity extraction accurate for the drugs clinicians actually prescribe. Efficiency and safety gains will be more durable if leadership invests in brief, role-specific training; routine process fidelity audits modeled on established reconciliation indicators; and lightweight feedback loops that capture acceptance, overrides, and reasons for dismissal to drive iterative tuning. From an informatics perspective, teams should establish governance for model stewardship: monitor extraction quality, track drift when documentation patterns or formularies change, and schedule periodic re-validation against pharmacist-curated reference sets. To extend impact beyond admission, discharge and post-discharge workflows should be aligned so that corrections propagate into the active medication list, patient instructions, and community pharmacy communication. Equity and usability deserve explicit attention; providing translated prompts, caregiver-friendly intake forms, and accessible interfaces can reduce missed medications among patients with language barriers or low health literacy. For operations, the observed reduction in completion time should be banked as capacity: health systems can expand pharmacist coverage hours, add targeted counseling for very high-risk patients, or reallocate time to complex reconciliations that still require deep clinical reasoning. Finally, to inform scale-up and reimbursement discussions, finance and quality teams should quantify avoided harm and downstream use changes, conduct sensitivity analyses for different admission volumes, and compare the cost of development and maintenance with recovered pharmacist time. In summary, pairing disciplined reconciliation practice with AI that assembles cleaner inputs and focuses expert attention offers a practical, scalable route to safer and faster admissions; with deliberate governance, thoughtful workflow integration, and continuous learning from real-world use, these gains can be sustained and extended across the hospitalization continuum.

References

1. World Health Organization. Standard Operating Protocol for Medication Reconciliation (High 5s Project). Geneva: WHO; 2014.
2. World Health Organization. High 5s Project: Medication Reconciliation Implementation Guide. Geneva: WHO; 2014.
3. The Joint Commission. National Patient Safety Goal NPSG.03.06.01: Maintain and communicate accurate patient medication information. Oakbrook Terrace (IL): TJC; 2024.
4. The Joint Commission. National Patient Safety Goals®: Medication reconciliation focus. Oakbrook Terrace (IL): TJC; 2025.
5. Barnsteiner J. Medication reconciliation. In: Hughes RG, editor. Patient Safety and Quality: An Evidence-Based Handbook for Nurses. Rockville (MD): AHRQ; 2008. p. 459-472.
6. Mueller SK, Sponsler KC, Kripalani S, Schnipper JL. Hospital-based medication reconciliation practices: a systematic review. Arch Intern Med. 2012;172(14):1057-1069.
7. Mueller SK, *et al*. A toolkit to disseminate best practices in inpatient medication reconciliation (MARQUIS). BMJ Qual Saf. 2013;22(12):968-979.
8. Kwan JL, Lo L, Sampson M, Shojania KG. Medication reconciliation during transitions of care: a systematic review. Ann Intern Med. 2013;158(5 Pt 2):397-403.
9. Mekonnen AB, McLachlan AJ, Brien JA. Effectiveness of pharmacist-led medication reconciliation programmes on clinical outcomes at hospital transitions: systematic review and meta-analysis. BMJ Open. 2016;6:e010003.
10. Gleason KM, McDaniel MR, Feinglass J, *et al*. Results of the Medications at Transitions and Clinical Handoffs (MATCH) study: admission medication reconciliation errors. J Gen Intern Med. 2010;25(5):441-447.
11. Quélennec B, Beretz L, Paya D, *et al*. Potential clinical impact of medication discrepancies at hospital admission in elderly patients. Eur J Intern Med. 2013;24(6):530-535.
12. Cornu P, Steurbaut S, Leysen T, *et al*. Effect of medication reconciliation at hospital admission on discrepancies during hospitalization and at discharge in geriatric patients. Ann Pharmacother. 2012;46(4):484-494.
13. Belda-Rustarazo S, Cantero-Hinojosa J, Salmerón-García A, *et al*. Medication reconciliation at admission

and discharge: an observational study. Int J Clin Pract. 2015;69(3):329-339.
14. Redmond P, Grimes TC, McDonnell R, *et al*. Impact of medication reconciliation for improving transitions of care. Cochrane Database Syst Rev. 2018;(8):CD010791.
15. Tamblyn R, *et al*. Effect of an electronic medication reconciliation intervention on adverse drug events and health care use: cluster randomized trial. JAMA Netw Open. 2019;2(9):e1910756.
16. Sardaneh AA, *et al*. Pharmacist-led admission medication reconciliation before and after an electronic medication management system. Int J Med Inform. 2017;102:1-8.
17. Xu H, Stenner SP, Doan S, *et al*. MedEx: a medication information extraction system for clinical narratives. J Am Med Inform Assoc. 2010;17(1):19-24.
18. Savova GK, Masanz JJ, Ogren PV, *et al*. Mayo cTAKES: architecture, component evaluation, and applications. J Am Med Inform Assoc. 2010;17(5):507-513.
19. Uzuner Ö, South BR, Shen S, DuVall SL. 2010 i2b2/VA challenge on concepts, assertions, and relations in clinical text. J Am Med Inform Assoc. 2011;18(5):552-556.
20. Jagannatha AN, Yu H. Structured extraction of medication entities and their attributes from clinical text using recurrent neural networks. J Am Med Inform Assoc. 2016;23(3):524-532.
21. Wei Q, Ji Z, Li Z, *et al*. A study of deep learning for medication and adverse event extraction from clinical text. JAMIA Open. 2019;2(3):379-384.
22. Kreimeyer K, Foster M, Pandey A, *et al*. Natural language processing systems for capturing and standardizing unstructured clinical information: a systematic review. J Biomed Inform. 2017;77:34-49.