

From Interaction Checking to Actionable Medication Safety

Designing Drug-Drug Interaction Decision Support for Physicians: An Evidence-Based Framework for the Hidoc Dr. Drug Interaction Capability

Dr Sadhana Gadia and Varun Gadia

Hidoc Journal | Evidence Review & Platform Evaluation Framework | 12 August 2026

Structured Abstract

Background: Drug-drug interactions (DDIs) are an important medication-safety concern, particularly in polypharmacy, multimorbidity, older adults, high-risk medicines and transitions of care. Digital interaction checkers can identify pharmacokinetic and pharmacodynamic risks quickly, but conventional alerting systems frequently generate warnings that clinicians override. Hidoc Dr. publicly describes a Drug Interaction capability that allows physicians to check potential interactions and adverse effects to support safer prescribing. The clinical value of such a capability depends not only on the completeness of an interaction database, but on severity prioritization, patient context, actionable management advice, evidence provenance, workflow design and continuous governance.

Objective: To synthesize current clinical, regulatory, informatics and medication-safety evidence relevant to digital DDI checking and to propose an evidence-based product, governance and prospective evaluation framework for Hidoc Dr.

Methods: A narrative evidence synthesis was undertaken using Hidoc Dr. public product information and authoritative sources including the World Health Organization (WHO) Medication Without Harm programme, the 2024 ICH M12/FDA Drug Interaction Studies guidance, systematic reviews of DDI alerts and clinical decision support, comparative studies of interaction knowledge bases, and peer-reviewed literature on alert fatigue, contextualized decision support and medication safety. Product-specific effectiveness was not inferred where independent validation data were unavailable.

Results: The evidence shows a persistent gap between interaction detection and clinically useful decision support. A 2024 systematic review and meta-analysis of 16 studies reported that physicians overrode approximately 90% of DDI alerts across 570,776 prescriptions. A 2025 systematic review of eight studies involving 43,413 patients found that DDI decision-support alerts modestly influenced prescribing behavior but did not significantly reduce targeted adverse interactions or improve patient-important outcomes. Knowledge bases also differ in interaction coverage and severity classification. These findings favor a patient-contextualized architecture that distinguishes mechanism, severity, probability, evidence strength and recommended action; reserves interruptive alerts for high-consequence situations; incorporates laboratory, renal, hepatic and demographic context where available; and makes evidence sources and override rationale auditable.

Conclusion: A drug-interaction platform should be designed as a medication-safety decision system rather than a binary pairwise checker. Hidoc Dr. can create greater clinical value by moving from detection toward prioritized, explainable and actionable interaction management. Prospective evaluation should measure computational and knowledge-base accuracy, clinically relevant sensitivity, alert burden, override appropriateness, prescribing changes, time burden, user trust and patient-important safety outcomes.

Keywords: drug-drug interactions; clinical decision support; medication safety; alert fatigue; polypharmacy; pharmacokinetics; prescribing; clinical informatics; adverse drug events; Hidoc Dr

1. Background

Medication prescribing increasingly occurs in patients with multiple chronic conditions, multiple prescribers and complex treatment regimens. The resulting combinations can be therapeutically necessary, intentionally synergistic, neutral, or harmful. A drug-drug interaction occurs when one medication modifies the effect or disposition of another. The mechanism may be pharmacokinetic - affecting absorption, distribution, metabolism or elimination - or

pharmacodynamic, where two agents produce additive, antagonistic or otherwise interacting physiologic effects. The clinical consequences range from trivial laboratory changes to treatment failure, severe toxicity, arrhythmia, bleeding, serotonin toxicity, hyperkalemia, respiratory depression or other serious outcomes.

WHO's Medication Without Harm initiative identifies medication safety as a global patient-safety priority and highlights polypharmacy, high-risk situations and transitions of care as key action areas [2,3]. WHO notes that unsafe medication practices and medication errors are leading causes of avoidable harm in healthcare and that medication error can arise throughout prescribing, transcribing, dispensing, administration and monitoring [2]. DDI management therefore belongs inside a wider medication-use system rather than being treated as an isolated database lookup.

Hidoc Dr. publicly lists Drug Interaction among its core physician offerings and describes the capability as allowing clinicians to check potential interactions and adverse effects to support safe prescribing practices [1]. This is a clinically relevant product position. A physician entering two or more medicines expects more than a theoretical statement that an interaction exists. The useful question is whether the interaction matters for this patient, how serious it is, what mechanism explains it, what evidence supports it, whether timing or dose alters risk, and what the clinician should do next.

The distinction between detection and decision support is central. Traditional interaction databases often operate pairwise: drug A plus drug B produces an alert. Yet clinical risk is conditional. The same pharmacologic interaction may be negligible in a young person with normal renal function and dangerous in an older patient with chronic kidney disease, baseline electrolyte abnormality, frailty or concurrent high-risk medicines. A medication-safety system that ignores this context may be technically sensitive but clinically noisy.

The consequences of excessive sensitivity are visible in the alert-fatigue literature. Clinicians repeatedly exposed to low-value warnings learn to dismiss them. A 2024 systematic review and meta-analysis by Felisberto and colleagues included 16 studies and found an overall physician DDI-alert override prevalence of approximately 90% across 570,776 prescriptions [4]. Such an override rate does not prove that clinicians are unsafe; it more plausibly reflects a system-level mismatch between the theoretical interaction universe and the set of warnings that actually require action at the moment of prescribing.

For Hidoc Dr., the opportunity is therefore to treat an interaction check as a structured medication-safety consultation. The product should detect relevant drug pairs, explain mechanisms, classify evidence and severity, identify patient-specific risk modifiers, provide management options and preserve clinician autonomy. The design target is not "maximum alerts." It is maximum clinically useful signal with minimal unnecessary interruption.

2. Product Analysis: What a Modern Hidoc Drug Interaction Checker Should Represent

A drug-interaction checker can be viewed as a multilayer knowledge system. The first layer is drug identity. Medicines may be entered by generic name, brand name, combination product, salt, formulation or class. Accurate normalization is essential because the system cannot reason correctly if the selected medicine is ambiguous. Brand-to-generic mapping must be market-aware, and combination products should expand into their active ingredients before interaction analysis.

The second layer is the interaction assertion. Rather than storing a simple yes/no relationship, a governed interaction record should include the implicated ingredients, mechanism, direction of effect, expected clinical consequence, evidence level, onset, severity, affected populations and recommended management. Pharmacokinetic interactions should identify the pathway where known - for example enzyme inhibition, induction or transporter effects - while pharmacodynamic interactions should identify the relevant physiologic mechanism such as QT prolongation, bleeding risk, central nervous system depression or potassium elevation.

The third layer is evidence provenance. The 2024 ICH M12 guidance harmonizes recommendations for evaluating enzyme- and transporter-mediated pharmacokinetic interaction potential during drug development [5]. Such regulatory evidence is valuable, but real-world interaction support must also integrate approved product labeling, clinical pharmacology studies, guidelines, safety communications, observational evidence and clinically accepted management

recommendations. The user should be able to see where an interaction statement came from and when it was last reviewed.

The fourth layer is clinical severity. Severity should not be confused with certainty. A rare but plausible interaction may have catastrophic potential, while a well-documented interaction may have only modest clinical consequences. Systems should therefore represent at least two axes: expected consequence and confidence/evidence. This prevents the common problem of collapsing "well established" and "clinically severe" into one label.

The fifth layer is patient context. Relevant modifiers include renal and hepatic function, age, pregnancy, baseline QT interval, potassium, magnesium, blood pressure, therapeutic drug levels, indication, dose, route, timing, treatment duration, pharmacogenomic factors and concomitant medicines. Not every context variable will be available in a standalone Hidoc application, but the architecture should be designed so that optional context can refine risk rather than requiring the system to behave identically for every patient.

The sixth layer is management. A warning without a management pathway creates cognitive burden. Management may include avoiding the combination, substituting one medicine, reducing dose, separating administration time, obtaining a baseline test, repeating a laboratory measurement, monitoring for a specific symptom, limiting duration or accepting the combination with documented clinical justification. The recommendation should indicate whether it is derived from a label, guideline, expert consensus or the platform's editorial interpretation.

The seventh layer is interaction hierarchy. A clinician entering six medicines creates 15 pairwise combinations before considering higher-order effects. Displaying 15 equal-priority warnings is not useful. The system should rank interactions by clinical consequence, patient vulnerability and actionability, then group related risks. Three medicines that prolong QT may represent a cumulative risk better communicated as one coherent warning than three separate pair alerts.

The eighth layer is lifecycle management. Drug labels change, new medicines enter markets, evidence emerges and safety communications alter recommendations. Every interaction assertion should therefore have a version, review date, source set and editorial owner. A DDI database is not a static encyclopedia; it is continuously maintained clinical knowledge.

3. The Evidence Problem: Interaction Databases Do Not Always Agree

Clinicians often assume that interaction databases provide equivalent answers. Comparative research shows otherwise. Fung and colleagues compared three commercial drug-interaction knowledge bases - First DataBank, Micromedex and Multum - and found meaningful differences in the drugs represented, interaction rules and agreement with a reference list of highly significant DDIs [6]. The study illustrates a fundamental challenge: the interaction knowledge base itself is an active clinical model, not a neutral index.

Variation can arise because vendors use different inclusion thresholds, terminology, evidence grading, severity categories and update processes. One database may classify an interaction as major, another as moderate and a third may not display it at all. A clinician who uses two different systems may therefore receive conflicting guidance without understanding why.

This variability matters for product governance. Hidoc should not claim completeness merely because its database contains a large number of interaction pairs. The clinically important question is whether the platform captures interactions that require meaningful management while avoiding unsupported or irrelevant assertions. A curated reference set of high-priority DDIs can serve as one validation layer, but no single external list represents the entire truth. Different specialties and patient populations encounter different hazards.

Knowledge-base evaluation should therefore be multidimensional. Coverage asks whether known important interactions are present. Precision asks how many displayed interactions are clinically meaningful in the intended context. Concordance examines whether mechanism, severity and management agree with authoritative sources. Currency assesses how quickly safety-relevant changes are incorporated. Traceability asks whether a reviewer can identify the evidence supporting each statement.

The challenge becomes larger with international deployment. Product labels, approved indications, formulations and regulatory warnings can vary by country. Brand names can differ dramatically.

Some medicines are unavailable or withdrawn in particular markets. A global Hidoc platform should separate universal pharmacology from market-specific regulatory content. The chemical interaction may be global; the recommended action and product availability may not be.

The platform should also distinguish established interactions from theoretical class effects. Mechanistic reasoning is important, particularly for new drugs, but theoretical interactions should not be presented with the same authority as well-characterized clinical effects. The evidence interface can state whether the assertion derives from human clinical studies, pharmacokinetic studies, case reports, in vitro findings, class inference or expert consensus.

Artificial intelligence may eventually assist evidence surveillance, but AI should not autonomously create live interaction rules. Natural-language systems can identify candidate safety updates, extract drug pairs from new labels or summarize new pharmacology, yet the final assertion should undergo clinical pharmacology and editorial review. The primary safety asset is not the language model; it is the governed knowledge base.

4. Alert Fatigue and Why More Warnings Can Produce Less Safety

Clinical decision support succeeds only when clinicians attend to it. DDI alerts have become a classic example of how excessive sensitivity can undermine that objective. The 2024 Felisberto systematic review screened 1,873 records and included 16 studies. Across the meta-analysis, DDI alerts were generated in approximately 13.7% of more than 21 million prescriptions, while the pooled override prevalence was 90% across 570,776 prescriptions [4]. Heterogeneity was extremely high, reflecting differences in systems and clinical settings, but the overall message is consistent with decades of informatics research: most DDI alerts are not acted upon.

Older studies reported similarly high override rates. Bryant and colleagues observed a drug-drug alert override rate of 95.1% in a hospital environment [7]. A 2022 scoping review by Villa Zapata and colleagues identified override rates ranging from approximately 55% to 98% across 34 studies [8]. Repetition of low-relevance alerts teaches clinicians that the expected value of reading the next alert is low.

Alert fatigue should not be interpreted simply as physician noncompliance. Alert acceptance depends on relevance, timing, severity, familiarity and availability of a practical alternative. If a warning fires after the clinician has already considered the interaction, or if the patient has tolerated the combination for years under monitoring, an interruptive alert adds little value. If the warning offers no actionable recommendation, the clinician must spend additional time reconstructing the problem.

The design response should be tiered interruption. High-consequence, time-sensitive interactions may justify an interruptive alert requiring acknowledgment or a reason for override. Moderate interactions may be shown non-interruptively within the medication review. Low-consequence or well-known combinations may be available on demand without interrupting workflow. AHRQ-funded work on meaningful drug-interaction alerts explicitly explores contextualization as a way to reduce alert burden while preserving important safety signals [9,10].

Alert presentation should also state why the warning applies now. "Potential interaction" is weaker than "Both medicines can increase serum potassium; this patient's last potassium was 5.4 mmol/L and eGFR is 28 mL/min/1.73 m²." The second alert connects pharmacology to patient vulnerability. Context increases actionability and can justify interruption.

Override itself should be treated as data. A high override rate for a particular interaction may indicate poor alert design, a common clinically acceptable combination, or a local practice pattern that deserves review. Conversely, frequent acceptance may signal that the rule is useful. Governance teams should review alert performance by interaction, setting and specialty rather than accepting the rule set as immutable.

Importantly, forcing extra authentication or complicated override steps does not solve low relevance. The 2024 meta-analysis cites evidence that burdensome override mechanisms can worsen frustration without necessarily improving safety [4]. Friction should be proportional to clinical risk.

5. Do DDI Alerts Improve Patient Outcomes?

The most important test of DDI decision support is not whether it generates alerts but whether it reduces meaningful harm. Evidence here is surprisingly limited.

Holbrook and colleagues published a 2025 systematic review evaluating electronic DDI alerts at prescribing, dispensing or medication administration [11]. The review included eight studies: seven randomized trials and one prospective cohort study, involving 43,413 patients. Four studies were conducted in Europe, three in North America and one in Israel, across inpatient and outpatient settings.

The findings were cautious. Mortality was similar between intervention and control groups. DDI alerts modestly influenced targeted prescribing behavior, with a pooled odds ratio of 2.08, but did not significantly reduce the incidence of targeted adverse interactions, for which the pooled odds ratio was 0.86 with a confidence interval crossing no effect [11]. The review concluded that current evidence does not show significant improvement in patient-important outcomes and emphasized the need for greater accuracy, clinical relevance and patient-specific integration.

Several explanations are plausible. First, many theoretical DDIs rarely produce clinical harm. Second, alerts may be delivered to patients with low baseline vulnerability. Third, clinicians may override warnings appropriately because the benefit of the combination outweighs the theoretical risk. Fourth, substituting a medicine after an alert may create a different adverse effect. Fifth, trials have often been underpowered for relatively rare serious outcomes.

This evidence should change how digital platforms define success. A Hidoc evaluation should not use "number of interactions detected" as the primary clinical outcome. Nor should alert acceptance automatically be interpreted as correct behavior. The relevant outcomes are whether clinically important interactions are recognized, whether management becomes more appropriate, whether preventable adverse events decrease and whether the system introduces new errors or unnecessary changes.

The review also highlights the need to study patient-important endpoints rather than only process measures. A system can increase laboratory testing and prescription changes without improving health outcomes. Process outcomes remain useful - especially during early product evaluation - but they should be linked to a plausible safety pathway.

For Hidoc, the safest scientific posture is therefore to describe the Drug Interaction tool as decision support, not as a proven mechanism for preventing adverse events unless prospective evidence demonstrates that effect. This distinction strengthens rather than weakens the product's credibility.

6. Patient-Contextualized Interaction Management

Patient-specific risk is the major frontier in DDI decision support. Many interactions are conditional. A potassium-raising combination becomes more concerning in renal impairment or with elevated baseline potassium. QT-prolonging combinations become more relevant in patients with electrolyte disturbance, bradycardia or pre-existing QT prolongation. Anticoagulant interactions depend on bleeding risk, dose, renal function, age and indication.

A contextualized Hidoc architecture can therefore operate in two modes. In basic mode, the clinician enters medicines only and receives conservative interaction information. In enhanced mode, the user can optionally enter relevant variables such as age, renal function, liver function, pregnancy, key laboratory values or indication. The platform then refines priority and management.

Context should suppress alerts only when suppression is evidence-based. The goal is not to reduce alert volume cosmetically. A suppressed alert should have a documented rule explaining why the interaction is not currently actionable. For example, an interaction that matters only with a prolonged treatment duration might be downgraded for a one-time dose. A route-specific interaction may be irrelevant when the medicine is administered topically.

Temporal context matters as well. Enzyme induction often develops and resolves over time; inhibition can have a different onset. Some interactions can be managed by separating administration times, whereas others cannot because the mechanism is systemic. The system should avoid generic advice such as "take two hours apart" unless timing truly mitigates the interaction.

Indication is another underused variable. Some combinations that appear interacting are deliberately prescribed together because their therapeutic benefit is well established. Dual antiplatelet therapy, combination antihypertensives or ritonavir-boosted regimens illustrate why "interaction" does not mean "contraindicated." The product should distinguish intentional pharmacologic combinations from potentially harmful unintended DDIs.

Patient context also supports monitoring recommendations. Rather than simply advising caution, the system can recommend a specific action: check potassium within a defined interval, monitor INR, obtain an ECG, observe for sedation, reduce dose or select an alternative. Such recommendations should be source-based and adaptable to local standards.

Where Hidoc does not have access to patient data, the interface can ask concise risk-refining questions only for high-priority interactions. This is preferable to requiring a full patient profile for every check. The system should request context when it can materially change the recommendation.

7. Regulatory and Clinical Pharmacology Foundations

The evidence underlying many interaction rules originates during drug development. The ICH M12 Drug Interaction Studies guideline, finalized in 2024 and adopted by FDA, provides harmonized recommendations for evaluating enzyme- and transporter-mediated pharmacokinetic DDIs [5]. It addresses in vitro assessment, clinical interaction studies and interpretation of interaction potential for investigational medicines. This guidance is directly relevant to the quality of upstream evidence incorporated into interaction databases.

However, a regulatory DDI study and a bedside decision are not the same thing. Regulatory studies often isolate mechanisms under controlled conditions, sometimes using healthy volunteers, probe substrates or strong inhibitors and inducers. Clinical decision support must translate those findings into management for heterogeneous patients taking multiple medicines.

Product labeling remains a critical source because it expresses regulator-reviewed interaction information and management recommendations. Yet labels can be conservative and may differ between jurisdictions. A platform should ideally distinguish label-based recommendations from independent clinical guidelines and other evidence.

The drug-interaction checker itself may also fall within software governance frameworks depending on intended use and functionality. A standalone reference that displays known interaction information is different from software that automatically recommends a personalized treatment change. As functionality becomes more patient-specific and directive, regulatory assessment may become more important.

Clinical pharmacology editorial oversight is therefore essential. A multidisciplinary team might include physicians, pharmacists, clinical pharmacologists, pharmacovigilance specialists and informaticians. High-risk changes should undergo dual review. Every production update should be traceable to a source and approved version.

The platform should also maintain a clear separation between scientific interaction content and pharmaceutical promotion. A sponsor should not be able to suppress a clinically relevant interaction involving its medicine or exaggerate an interaction involving a competitor. Sponsored educational material should be visually separate from the interaction knowledge engine.

For international use, the governance framework should map medicines to jurisdiction-specific labels and safety communications while retaining a common scientific core. This allows Hidoc to scale globally without pretending that one regulatory document is universal.

8. International Use Cases

In primary care, DDI checking is especially relevant to older adults with multimorbidity. A physician may inherit prescriptions from cardiology, psychiatry, endocrinology and other specialties. The useful interaction tool is one that prioritizes combinations likely to matter and supports medication review rather than simply returning dozens of warnings.

In hospitals, DDI support can be embedded into prescribing and medication reconciliation. Inpatient context offers richer data, including current laboratory values and vital signs, but also greater alert burden. High-frequency orders and protocolized combinations make contextual suppression particularly important.

In pharmacy practice, interaction checking can support dispensing decisions, patient counseling and escalation to the prescriber. The management language should therefore distinguish what a pharmacist can resolve through timing or counseling from what requires prescription change.

In oncology, transplant medicine, infectious diseases and other specialties with complex pharmacology, interaction management may require therapeutic drug monitoring, dose adjustment or temporary interruption. Ritonavir-containing regimens, calcineurin inhibitors, oral anticancer therapies and anticoagulants illustrate contexts where interaction advice must be precise and current.

In resource-limited settings, a mobile interaction tool may substitute for access to a clinical pharmacist or expensive enterprise knowledge base. This increases the value of portability but also increases responsibility for evidence quality. Offline caching should preserve the interaction database version and prompt synchronization when updates become available.

For medical education, Hidoc can use interaction checks as teaching objects. Rather than only displaying the final warning, an educational mode can explain metabolic pathways, transporter effects, pharmacodynamic mechanisms and why patient context changes risk. This aligns the platform's medication-safety function with continuing professional development.

For transitions of care, medication lists are often incomplete or changed rapidly. Interaction checking can be paired with medication reconciliation, but the system should explicitly display whether the list is patient-reported, discharge-derived or verified. Accurate input is a prerequisite for accurate interaction analysis.

Across all of these use cases, the core principle remains stable: the interaction engine should help the clinician decide what to do, not merely prove that two drugs have been mentioned together in pharmacology literature.

9. Implementation Architecture for Hidoc Dr.

A practical implementation can begin with a normalized medication layer. Drug names should resolve to standardized active ingredients, strengths, routes and combination products. Duplicate ingredients should be detected before interaction analysis because therapeutic duplication can itself create risk.

The next component is the interaction rules engine. Each rule should contain a unique identifier, drug pair or class relationship, mechanism, severity, evidence grade, onset, management, monitoring, patient modifiers, source references and version. Rules involving classes should specify inheritance logic so that newly added medicines do not automatically acquire unsupported class interactions.

A context engine can then evaluate optional patient variables. Rules should declare which variables matter. This avoids collecting unnecessary data. If renal function is irrelevant to a particular interaction, it should not influence priority. If a specific laboratory threshold materially changes risk, the threshold should be explicit and sourced.

The presentation layer should separate summary from detail. The initial result can show a ranked list: avoid, major/action required, monitor/adjust, minor/informational, and no known clinically meaningful interaction. Expanding a result reveals mechanism, evidence, management alternatives and references.

Interruptive alerts should be reserved for the highest-priority scenarios and used primarily when Hidoc is embedded in an active prescribing workflow. A standalone checker can present severity without forcing acknowledgment because the user has intentionally requested the information.

Search speed and mobile usability matter, but they are subordinate to interpretability. The clinician should never need to decode unexplained icons or vendor-specific severity terminology. Labels should be defined in plain clinical language.

A quality dashboard should monitor rule age, evidence updates, reported errors, disagreement with external references, frequently viewed interactions, user-reported clinical irrelevance and override data when available. High-use and high-severity rules deserve shorter review intervals.

Pharmacovigilance links can create an additional safety pathway. If a clinician encounters an unexpected interaction-related adverse event,

Hidoc can guide the user toward appropriate reporting channels without conflating spontaneous reporting with the interaction database itself.

Finally, the platform should support a formal correction mechanism. Users should be able to report a missing interaction, incorrect severity, outdated management recommendation or wrong brand mapping. Reports should be triaged according to potential clinical harm and produce a visible version update when confirmed.

10. Governance, Safety and Limitations

No interaction database can eliminate uncertainty. Even highly established pharmacokinetic changes do not translate into identical clinical outcomes for every patient. Conversely, absence of a listed interaction does not prove absence of risk, especially for new medicines, supplements, traditional medicines or poorly studied combinations.

The platform should therefore avoid absolute language unless the source supports it. "No known clinically significant interaction in the reviewed sources" is more defensible than "safe combination." Similarly, "avoid" should be reserved for interactions with a strong management basis rather than used as a generic severity label.

Supplements and herbal products require special caution because composition, dose and evidence quality can vary. Hidoc can include known interaction information while clearly differentiating robust evidence from case reports or theoretical mechanisms.

Pregnancy, pediatrics and severe organ dysfunction are additional high-risk contexts. Interaction evidence may be sparse, and dose consequences may differ. The interface should surface uncertainty and recommend specialist or pharmacist review where appropriate rather than over-automating.

Privacy becomes relevant when patient variables are entered. A standalone checker can often process data transiently without storing identifiable information. If Hidoc saves medication profiles, synchronization or patient identifiers, stronger access control, encryption, retention and consent governance are required.

AI-generated explanations must be constrained. Generative AI may improve readability or answer follow-up questions, but it should retrieve from the governed interaction record rather than inventing novel interaction claims. The deterministic knowledge base should remain the source of truth. AI output should be reproducible enough for audit in high-risk contexts.

Commercial governance should ensure that interaction severity and management are insulated from advertiser influence. The editorial function must have authority to publish safety information even when commercially inconvenient.

A limitation of this article is that it is a narrative synthesis rather than a systematic review. A second limitation is that public Hidoc information establishes the existence and intended purpose of its Drug Interaction capability but does not provide independent validation data on database coverage, computational accuracy, alert performance or patient outcomes. No such product-specific claims are made.

A third limitation is the evidence base itself. The 2025 Holbrook systematic review found only eight eligible higher-quality studies and judged most to have high risk of bias [11]. The absence of demonstrated patient benefit does not prove that well-designed DDI support is ineffective; it demonstrates that current evidence is limited and that future systems should be evaluated more rigorously.

Finally, DDI checking is only one component of medication safety. Dose appropriateness, contraindications, allergies, therapeutic duplication, renal adjustment, adherence, monitoring and medication reconciliation all interact with DDI risk. The platform should avoid creating a false sense of safety from a clean interaction screen.

11. Proposed Evaluation Methodology for Hidoc Dr.

A prospective evaluation program should begin with knowledge-base validation rather than user engagement metrics.

Phase 1 should establish a reference set of clinically important interactions across major specialties. The set should include pharmacokinetic and pharmacodynamic interactions, contraindicated combinations, interactions requiring monitoring, route-specific interactions, dose-dependent interactions and deliberately noninteracting controls. An independent expert panel should define expected classification and management before testing the platform.

Phase 2 should measure rule-level accuracy. Outcomes include sensitivity for reference interactions, false-positive rate for negative controls, concordance of severity, concordance of mechanism and adequacy of management advice. Errors should be weighted by potential clinical consequence rather than pooled into one percentage.

Phase 3 should perform knowledge-base comparison against at least two established external sources. Disagreements should undergo blinded clinical review rather than assuming the majority is correct. This process can identify both Hidoc defects and inconsistencies in external databases.

Phase 4 should evaluate patient-context refinement using synthetic clinical scenarios. Each drug pair can be tested across changes in renal function, electrolytes, age, dose, route and duration. The primary outcome is whether the system appropriately changes priority or management when the contextual variable is clinically relevant and remains stable when it is irrelevant.

Phase 5 should test usability with practicing physicians and pharmacists. Participants should complete standardized prescribing scenarios. Outcomes include time to identify the important interaction, correct interpretation, management choice, unnecessary medication changes, navigation errors and perceived cognitive load. Think-aloud interviews can reveal why an apparently correct alert is ignored.

Phase 6 should evaluate alert burden in a simulated prescribing workflow. The same medication orders can be processed through a conventional noncontextual rule set and the Hidoc contextualized system. Outcomes include alerts per 100 orders, high-priority alert sensitivity, proportion of interruptive alerts judged actionable and time burden.

Phase 7 should evaluate override appropriateness. In a live or controlled implementation, accepted and overridden high-priority alerts should be reviewed by an independent panel. The desired endpoint is not the lowest override rate; it is the highest proportion of clinically appropriate responses.

Phase 8 should conduct a prospective comparative clinical study. Where feasible, practices or wards could be randomized or stepped into Hidoc-supported DDI checking versus usual care. Process outcomes could include inappropriate interacting prescriptions, monitoring completion and clinically justified medication changes. Patient-important outcomes should include interaction-related adverse events, emergency visits, hospitalization and other prespecified safety events, recognizing that large sample sizes may be required.

Phase 9 should measure clinician burden and trust. The 2025 systematic review noted that burden outcomes specific to DDI alerts were rarely measured [11]. Hidoc should therefore prospectively capture time per alert, perceived interruption, trust, alert fatigue, satisfaction and willingness to use the tool again.

Phase 10 should test update governance. Simulated new regulatory warnings should be introduced into a staging environment. The evaluation should measure detection, evidence review, rule creation, independent approval, regression testing and release time. A clinically accurate but slowly updated interaction database may still be unsafe.

All studies should prespecify outcomes and distinguish exploratory from confirmatory analyses. Product versions must be recorded. If the rule set changes during a study, the version and date should be traceable. Negative findings should be reported alongside positive findings.

The ultimate evidence standard should be layered. Computational and editorial validation can establish that Hidoc accurately represents its knowledge base. Usability studies can establish that clinicians understand it. Workflow studies can establish that it reduces low-value alerts. Only adequately designed clinical studies can establish whether it improves patient-important medication-safety outcomes.

12. Conclusion

Drug-drug interaction checking is a necessary but deceptively difficult component of medication safety. The pharmacologic literature contains thousands of possible interactions, yet clinicians need a much smaller set of timely, credible and actionable signals for the patient in front of them.

Hidoc Dr. publicly positions Drug Interaction as a core physician tool for identifying potential interactions and adverse effects [1]. The evidence reviewed here suggests that the highest-value evolution of that capability is not simply expansion of the interaction catalogue. It is contextualization.

The 2024 meta-analysis showing approximately 90% DDI-alert overrides demonstrates the limitations of indiscriminate warning systems [4]. The 2025 systematic review involving 43,413 patients further shows that conventional DDI alerts may influence prescribing without clear evidence of improvement in patient-important outcomes [11]. These findings should not lead to abandonment of DDI decision support. They should lead to better decision support.

A modern Hidoc interaction system should combine a curated, versioned knowledge base with evidence provenance, patient-specific modifiers, severity and confidence separation, clear management, restrained interruption, user feedback and independent editorial governance. Its performance should be evaluated not by alert volume but by clinically important sensitivity, actionable precision, appropriate clinician response, burden and ultimately patient safety.

The central product question should change from "Do these drugs interact?" to "Does this interaction matter for this patient now, and what should the clinician do about it?" A system that can answer that question reliably would move Hidoc Dr. from interaction checking toward actionable medication-safety infrastructure.

Declarations

Funding: No external funding is reported for this narrative evidence review.

Conflicts of interest: The authors are associated with Hidoc. Public product descriptions are separated from independent evidence, and no unvalidated Hidoc performance claims are made.

Data availability: No primary participant dataset was generated.

Ethics: No human participants were enrolled. This manuscript is a narrative evidence synthesis and proposed evaluation framework.

References

1. Hidoc Dr. Medical Learning Platform for Doctors. Core Offering: Drug Interaction. Hidoc Dr.; accessed 12 August 2026.
2. World Health Organization. Medication Without Harm. WHO Global Patient Safety Challenge. Accessed 12 August 2026.
3. World Health Organization. Medication safety in polypharmacy: technical report. Geneva: WHO; 2019.
4. Felisberto M, Lima GDS, Celuppi IC, et al. Override rate of drug-drug interaction alerts in clinical decision support systems: a brief systematic review and meta-analysis. *Health Informatics J.* 2024;30:14604582241263242. doi:10.1177/14604582241263242.
5. International Council for Harmonisation; US Food and Drug Administration. M12 Drug Interaction Studies. Final Guidance. August 2024.
6. Fung KW, Kapusnik-Uner J, Cunningham J, Higby-Baker S, Bodenreider O. Comparison of three commercial knowledge bases for detection of drug-drug interactions in clinical decision support. *J Am Med Inform Assoc.* 2017;24(4):806-812. doi:10.1093/jamia/ocx010.
7. Bryant AD, Fletcher GS, Payne TH. Drug interaction alert override rates in the Meaningful Use era. *Appl Clin Inform.* 2014;5(3):802-813.
8. Villa Zapata L, Subbian V, Boyce RD, et al. Overriding drug-drug interaction alerts in clinical decision support systems: a scoping review. *Stud Health Technol Inform.* 2022;290:380-384. doi:10.3233/SHTI220101.
9. Agency for Healthcare Research and Quality. Meaningful Drug Interaction Alerts. Digital Healthcare Research.
10. Agency for Healthcare Research and Quality. Creating Meaningful Decision Support to Reduce Drug-Drug Interaction Alerts. Digital Healthcare Research.
11. Holbrook AM, Matos Silva J, Faruque JAY, et al. Effect of electronic drug-drug interaction alerts on patient and clinician outcomes: a systematic review. *J Am Med Inform Assoc.* 2025;32(10):1617-1627.
12. Dechanont S, Maphanta S, Butthum B, Kongkaew C. Hospital admissions/visits associated with drug-drug interactions: a systematic review and meta-analysis. *Pharmacoepidemiol Drug Saf.* 2014;23(5):489-497.
13. Zheng WY, Richardson LC, Li L, Day RO, Westbrook JI, Baysari MT. Drug-drug interactions and their harmful effects in hospitalised patients: a systematic review and meta-analysis. *Eur J Clin Pharmacol.* 2018;74:15-27.

14. Fitzmaurice MG, Wong A, Akerberg H, et al. Evaluation of potential drug-drug interactions in adults in the intensive care unit: a systematic review and meta-analysis. *Drug Saf.* 2019;42:1035-1044.
15. Nabovati E, Vakili-Arki H, Taherzadeh Z, et al. Information technology-based interventions to improve drug-drug interaction outcomes: a systematic review on features and effects. *J Med Syst.* 2017;41:12.
16. Wong K, Yu SK, Holbrook A. A systematic review of medication safety outcomes related to drug interaction software. *J Popul Ther Clin Pharmacol.* 2010;17:e243-e255.
17. McGreevey JD 3rd, Mallozzi CP, Perkins RM, Shelov E, Schreiber R. Reducing alert burden in electronic health records: state of the art recommendations from four health systems. *Appl Clin Inform.* 2020;11:1-12.
18. Olakotan OO, Mohd Yusof M. The appropriateness of clinical decision support systems alerts in supporting clinical workflows: a systematic review. *Health Informatics J.* 2021;27:14604582211007536.
19. Muylle KM, Gentens K, Dupont AG, Cornu P. Evaluation of an optimized context-aware clinical decision support system for drug-drug interaction screening. *Int J Med Inform.* 2021;148:104393.
20. Chou E, et al. Designing and evaluating contextualized drug-drug interaction algorithms for clinical decision support. *J Am Med Inform Assoc.* 2021.
21. Van de Sijpe G, Quintens C, Walgraeve K, et al. Overall performance of a drug-drug interaction clinical decision support system: quantitative evaluation and end-user survey. *BMC Med Inform Decis Mak.* 2022.
22. Payne TH, Hines LE, Chan RC, et al. Recommendations to improve the usability of drug-drug interaction clinical decision support alerts. *J Am Med Inform Assoc.* 2015;22(6):1243-1250.
23. Phansalkar S, Desai AA, Bell D, et al. High-priority drug-drug interactions for use in electronic health records. *J Am Med Inform Assoc.* 2012;19(5):735-743.
24. Tilson H, Hines LE, McEvoy G, et al. Recommendations for selecting drug-drug interactions for clinical decision support. *Am J Health Syst Pharm.* 2016;73(8):576-585.
25. Kesselheim AS, Cresswell K, Phansalkar S, Bates DW, Sheikh A. Clinical decision support systems could be modified to reduce alert fatigue while still minimizing the risk of litigation. *Health Aff.* 2011;30(12):2310-2317.
26. Ancker JS, Edwards A, Nosal S, et al. Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system. *BMC Med Inform Decis Mak.* 2017;17:36.
27. Nanji KC, Seger DL, Slight SP, et al. Medication-related clinical decision support alert overrides in inpatients. *J Am Med Inform Assoc.* 2018;25(5):476-481.
28. Lapane KL, Waring ME, Schneider KL, Dube C, Quilliam BJ. A mixed method study of the merits of e-prescribing drug alerts in primary care. *J Gen Intern Med.* 2008;23:442-446.
29. van der Sijs H, Aarts J, Vulto A, Berg M. Overriding of drug safety alerts in computerized physician order entry. *J Am Med Inform Assoc.* 2006;13(2):138-147.
30. Weingart SN, Toth M, Sands DZ, Aronson MD, Davis RB, Phillips RS. Physicians' decisions to override computerized drug alerts in primary care. *Arch Intern Med.* 2003;163:2625-2631.
31. Shahmoradi L, et al. Clinical decision support systems-based interventions to improve medication outcomes: a systematic literature review. *Int J Med Inform.* 2021.
32. Kwan JL, Lo L, Ferguson J, et al. Computerised clinical decision support systems and absolute improvements in care: meta-analysis of controlled clinical trials. *BMJ.* 2020;370:m3216.
33. World Health Organization. Medication without harm: policy brief. Geneva: WHO; 2024.
34. World Health Organization. Global burden of preventable medication-related harm in health care: a systematic review. Geneva: WHO; 2024.
35. European Medicines Agency. Guideline on the investigation of drug interactions. EMA clinical pharmacology guidance and ICH M12 implementation resources.
36. US Food and Drug Administration. Drug Development and Drug Interactions: Table of Substrates, Inhibitors and Inducers. FDA.
37. Hines LE, Murphy JE. Potentially harmful drug-drug interactions in the elderly: a review. *Am J Geriatr Pharmacother.* 2011;9(6):364-377.
38. Malone DC, Abarca J, Hansten PD, et al. Identification of serious drug-drug interactions: results of the partnership to prevent drug-drug interactions. *J Am Pharm Assoc.* 2004;44:142-151.
39. Hines LE, Saverno KR, Warholak TL, et al. Pharmacists' and physicians' opinions on improving drug-drug interaction alerts. *J Am Med Inform Assoc.* 2012;19(6):941-945.
40. Wright A, Aaron S, Seger DL, Samal L, Schiff GD, Bates DW. Reduced effectiveness of interruptive drug-drug interaction alerts after repeated exposure. *J Am Med Inform Assoc.* 2019.
41. Bell BG, et al. The effect of computerized alerts on prescribing and patient outcomes: systematic review. 2025.
42. Baysari MT, et al. Effectiveness of computerised alerts to reduce drug-drug interaction exposure and alert burden. *BMJ Qual Saf.* 2025.