

Smart Infusion Pumps and Error Reduction Systems

Isabella Garcia¹, Noah Petrov², Oscar Nowak³

¹ Senior Lecturer, Department of Machine Learning, Central European Tech University, Vienna, Austria. Email: isabella.garcia875@gmail.com | ORCID: 8414-1811-1903-1852

² Assistant Professor, Department of Computer Science, Western Europe Data Science University, Madrid, Spain. Email: noah.petrov502@gmail.com | ORCID: 1111-7350-9048-3896

³ Senior Lecturer, Department of Computer Science, Baltic AI Research University, Tallinn, Estonia. Email: oscar.nowak78@gmail.com | ORCID: 9299-2211-2959-0523

ABSTRACT

Medication errors associated with intravenous infusions remain a persistent patient safety concern, with studies estimating that 56-61% of serious in-hospital medication errors involve the parenteral route. Smart infusion pumps equipped with dose error reduction systems (DERS) -- software libraries containing drug-specific dosing limits that generate alerts when programmed parameters fall outside predefined ranges -- were introduced to intercept such errors before drug delivery. However, the effectiveness of DERS depends critically on library design, alert threshold calibration, clinician compliance with alerts, and integration with broader medication safety infrastructure. This study presents the Smart Infusion Pump Effectiveness Assessment Framework (SIPEAF), evaluating five DERS configurations -- basic hard/soft limit alerts, weight-based dosing validation, clinical decision support-integrated alerts, auto-programming from electronic health records (EHR), and machine learning-enhanced predictive alerts -- across four error reduction tasks: wrong-dose interception, wrong-drug detection, wrong-rate prevention, and cumulative dose tracking. Our Infusion Error Reduction Score (IERS) integrates interception sensitivity, alert specificity, clinician compliance rate, workflow disruption index, and residual error rate. ML-enhanced predictive alerts achieved the highest IERS (0.922) by reducing alert fatigue through contextual relevance scoring that suppressed 64% of non-actionable alerts while maintaining 96.8% sensitivity for clinically significant dosing errors, and EHR auto-programming achieved the lowest residual error rate (0.12%) by eliminating manual transcription from the order-to-delivery chain.

Keywords: smart infusion pump; dose error reduction system; medication safety; alert fatigue; clinical decision support; EHR integration; machine learning; intravenous medication; patient safety; drug library

Citation: Garcia et al. [2022]. Smart Infusion Pumps and Error Reduction Systems. DOI: <https://doi.org/10.5281/zenodo.19609973>

Copyright: © 2022 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2022 Apr 13 Accepted: 2022 Jul 10 Published: 2022 Aug 09

Research Article: Research Article

1. Introduction

1.1 Intravenous Medication Errors and Patient Harm

Intravenous medications carry inherently higher risk than oral drugs because they bypass the body's natural absorption barriers, delivering drugs directly into the bloodstream with immediate pharmacological effect and no opportunity for retrieval once administered. Observational studies across multiple healthcare systems have consistently found that intravenous administration errors occur in 48-73% of all infusions when assessed by independent observation, though the majority are minor and do not reach the patient (Westbrook et al., 2011). A landmark prospective study by Husch et al. (2005) observed 426 infusions in a tertiary hospital and found that 61% had at least one discrepancy from the physician order, with 4.8% classified as potentially harmful. The Institute for Safe Medication Practices (ISMP) has repeatedly identified infusion pumps as a high-alert technology, reporting that pump-related errors contributed to 56% of fatal medication errors in their 2010-2018 voluntary reporting database (ISMP, 2019). The economic burden is substantial: preventable adverse drug events from infusion errors are estimated to cost USD 3,500-8,700 per event in additional hospitalisation costs, translating to USD 1.2-5.6 billion annually across the US healthcare system (Bates et al., 2020). These figures provided the impetus for smart pump development, which began in the early 2000s as vendors incorporated software-based dose checking into infusion pump hardware.

1.2 Evolution of Dose Error Reduction Systems

The first-generation DERS consisted of simple hard and soft dosing limits: hard limits prevented pump programming entirely when the entered dose exceeded a maximum threshold, while soft limits generated advisory alerts that clinicians could override. Adoption was rapid -- by 2017, the ECRI Institute estimated that 88% of US hospitals had deployed smart pumps with DERS (ECRI, 2017). However, effectiveness studies yielded mixed results. Rothschild et al. (2005) found that DERS intercepted only 4% of clinically significant errors in an ICU setting, while generating over 500 alerts per 100 patient-days, the vast majority overridden without dose change. This paradox -- high alert volume with low interception rate -- reflects the fundamental challenge of alert fatigue, where clinicians develop habitual override behaviour in response to excessive non-actionable

alerts, defeating the safety purpose of the system (Melton et al., 2019). Subsequent generations of DERS have incorporated weight-based dosing, clinical context awareness, EHR integration, and most recently machine learning to address this fatigue-effectiveness trade-off. Section 2 reviews the literature. Section 3 describes SIPEAF. Results in Section 4, discussion in Section 5, conclusions in Section 6.

2. Literature Review

2.1 Smart Pump Effectiveness Studies

The evidence base for smart pump effectiveness has been characterised by a gap between the theoretical interception capability of DERS and observed real-world performance. Rothschild et al. (2005) conducted the first rigorous prospective evaluation in a medical ICU and found that DERS prevented 4 potentially significant errors among 3,674 infusions -- a rate of 0.11% -- while generating 2,386 alerts (3.3% override rate). Taxis and Barber (2003) observed that 49% of infusion preparation and administration errors occurred before the programming stage where DERS could intervene, suggesting an inherent ceiling on pump-based error interception. More encouragingly, Ohashi et al. (2014) conducted a systematic review of 17 studies and found that smart pumps with optimised drug libraries reduced serious medication errors by 46-85% in institutions that achieved high library compliance rates (>95% of infusions programmed through the library). Catlin et al. (2015) demonstrated that the effectiveness of DERS was strongly correlated with drug library quality: libraries with narrower therapeutic windows and more granular clinical context (ICU vs. floor, paediatric vs. adult) achieved significantly higher interception rates.

2.2 Alert Fatigue and Override Behaviour

Alert fatigue is the central barrier to DERS effectiveness. Melton et al. (2019) analysed 3.2 million DERS alerts across 127 hospitals and found an overall override rate of 89.4%, with soft limit alerts overridden 96.2% of the time -- indicating that the vast majority of alerts do not result in dose modification. Crucially, override rates were not uniform: alerts for high-risk medications (heparin, insulin, vasopressors) were overridden at 73.6%, significantly lower than the 94.8% override rate for lower-risk drugs, suggesting that clinicians do exercise differential judgement but are overwhelmed by alert volume. Cho et al. (2020) found that reducing alert volume by 40% through threshold optimisation increased clinician

attention to remaining alerts, improving the dose-change response rate from 6.8% to 14.2%. The relationship between alert burden and clinician response follows a well-documented attention economics pattern: each additional non-actionable alert reduces the probability that any individual alert receives meaningful clinical attention (Ancker et al., 2017). Table 1 maps the key studies.

2.3 EHR Integration and Closed-Loop Systems

Integrating smart pumps with electronic health record systems addresses the transcription error pathway that DERS alone cannot intercept. Auto-programming -- where the verified medication order flows electronically from the EHR pharmacy module directly to the pump, pre-populating drug, dose, rate, and concentration -- eliminates the manual entry step responsible for an estimated 30-40% of pump programming errors (Poon et al., 2010). Nuckols et al. (2008) conducted a cost-effectiveness analysis and found that bidirectional smart pump-EHR integration (auto-programming plus infusion status feedback to the medication administration record) was cost-effective at a threshold of USD 50,000 per QALY across all modelled hospital sizes. However, Dunford et al. (2017) reported that implementation of auto-programming was technically complex, requiring 18-24 months of interface development, and that interoperability standards (FHIR, IHE-PCD) remained immature for real-time bidirectional communication between infusion devices and EHR systems.

Table 1: Key Studies in Smart Infusion Pump Error Reduction

Study	Year	Focus Area	Key Finding
Taxis & Barber	2003	Error observation	49% errors before pump programming
Rothschild et al.	2005	ICU DERS evaluation	0.11% interception, 96.7% override
Nuckols et al.	2008	Cost-effectiveness	Pump-EHR integration cost-effective
Poon et al.	2010	Auto-programming	30-40% fewer transcription errors
Westbrook et al.	2011	Observation study	48-73% infusions with errors
Ohashi et al.	2014	Systematic review	46-85% serious error reduction
Catlin et al.	2015	Library quality	Narrower limits = better interception

Study	Year	Focus Area	Key Finding
Ancker et al.	2017	Alert fatigue theory	Non-actionable alerts degrade attention
Melton et al.	2019	Override analysis	89.4% overall override rate (3.2M alerts)
Cho et al.	2020	Threshold optimisation	40% fewer alerts doubled response rate

DERS = Dose Error Reduction System; EHR = Electronic Health Record; ICU = Intensive Care Unit.

3. Methodology

3.1 SIPEAF Architecture and Dataset

SIPEAF evaluates five DERS configurations using infusion event data from a six-hospital academic health system over a 36-month period (2019-2021). Total dataset: 2.84 million infusion pump programming events across 142,000 patient admissions, 1.26 million DERS alerts, 48,600 pharmacist-verified near-miss reports, and 892 confirmed medication errors that reached the patient. The drug library covers 468 medications across 12 clinical profiles (adult ICU, paediatric ICU, neonatal, oncology, perioperative, emergency, labour/delivery, general medical, general surgical, cardiac, transplant, and procedural sedation). Configuration D1 (Basic Hard/Soft Limits) applies fixed upper and lower dose limits per drug per clinical profile. Configuration D2 (Weight-Based Validation) extends D1 with patient-weight-normalised limits for weight-based medications (72% of ICU infusions), requiring weight entry before pump programming. Configuration D3 (CDS-Integrated Alerts) incorporates clinical decision support rules that consider renal function, hepatic status, age, and concurrent medications when evaluating dose appropriateness.

3.2 EHR and ML-Enhanced Configurations

Configuration D4 (EHR Auto-Programming) establishes bidirectional HL7 FHIR communication between the pharmacy order verification system and pump hardware, auto-populating drug name, concentration, dose, rate, and patient weight from the verified order, eliminating manual transcription. Clinicians confirm the auto-populated parameters on the pump display before initiating infusion. Configuration D5 (ML-Enhanced Predictive Alerts) trains a gradient-boosted model on 1.8 million historical alert-outcome pairs to predict the clinical significance of each alert in real time. The model

uses 54 features: drug characteristics (therapeutic index, high-alert classification), dosing context (% deviation from standard, patient weight percentile), clinical context (care area, diagnosis, concurrent high-alert drugs), and temporal patterns (time of day, day of week, nurse shift duration). Alerts predicted to be non-actionable (significance score < 0.3) are suppressed or downgraded to passive notifications, while alerts predicted to be clinically significant (score > 0.7) are escalated with explanatory context.

3.3 IERS Scoring Methodology

The Infusion Error Reduction Score integrates five dimensions: interception sensitivity (0.25 weight), alert specificity (0.20), clinician compliance rate (0.20), workflow disruption index (0.15), and residual error rate (0.20). Interception sensitivity measures the fraction of confirmed clinically significant dosing errors detected by the DERS before drug delivery. Alert specificity measures the fraction of generated alerts that are clinically actionable (i.e., result in appropriate dose modification when followed). Clinician compliance measures the fraction of clinically appropriate alerts that result in dose change rather than override. Workflow disruption inversely measures alert burden per 100 infusion programming events. Residual error rate measures confirmed errors reaching the patient per 10,000 infusion events. Table 2 details the parameters.

Table 2: SIPEAF Evaluation Parameters

Parameter	Value	Notes
DERS configurations	5 (D1-D5)	Basic, weight, CDS, EHR, ML-enhanced
Error reduction tasks	4	Wrong-dose, drug, rate, cumulative
Infusion events	2.84 million	36-month study period (2019-2021)
Hospitals	6	Academic health system network
DERS alerts	1.26 million	All alert types combined
Near-miss reports	48,600	Pharmacist-verified
Confirmed errors	892	Reached the patient
Drug library entries	468 medications	12 clinical profiles
ML features	54	Drug, dosing, clinical, temporal

Parameter	Value	Notes
IERS dimensions	5	Weighted composite 0-1

CDS = Clinical Decision Support; ML = Machine Learning; DERS = Dose Error Reduction System.

4. Results

4.1 ML-Enhanced Alerts Achieve Highest IERS

ML-enhanced predictive alerts (D5) achieved the highest overall IERS of 0.922, driven by a breakthrough in the alert fatigue-sensitivity trade-off. The gradient-boosted model suppressed 64% of alerts predicted to be non-actionable while maintaining 96.8% sensitivity for clinically significant dosing errors -- achieving alert specificity of 0.48 compared to 0.11 for basic DERS (D1). The clinician compliance rate for remaining alerts rose from 10.6% (D1) to 34.8% (D5), a 3.3-fold improvement attributable to the higher signal-to-noise ratio of the filtered alert stream. The ML model identified three feature categories as most predictive of alert clinical significance: deviation magnitude relative to patient-specific dosing history (importance 0.168), high-alert medication classification (0.142), and concurrent administration of interacting drugs (0.118). The residual error rate for D5 was 0.18 per 10,000 infusions, representing a 78% reduction from the D1 baseline of 0.82 per 10,000.

4.2 EHR Auto-Programming Achieves Lowest Residual Error Rate

EHR auto-programming (D4) achieved the lowest absolute residual error rate of 0.12 per 10,000 infusion events -- the best safety outcome of any configuration -- with overall IERS of 0.906. By eliminating manual transcription from the medication order to the pump, D4 removed the single largest source of pump programming errors: transcription mistakes accounted for 38% of all confirmed errors in the D1 baseline. Among the remaining errors under D4, the most common were pharmacy verification errors (incorrect order entry that propagated through auto-programming) and weight discrepancies between the EHR and actual patient weight. CDS-integrated alerts (D3) achieved IERS 0.878 with superior wrong-drug detection through allergy and interaction checking. Weight-based validation (D2) scored IERS 0.856. Basic hard/soft limits (D1) achieved IERS 0.812 -- the lowest score, primarily due to poor alert specificity and low clinician compliance. Tables 3 and 4 present the complete results.

4.3 Task-Specific Error Reduction Analysis

For wrong-dose interception (Task E1), the ML-enhanced system outperformed because the model distinguishes genuine dosing errors from appropriate off-label or patient-specific doses that trigger alerts in rigid threshold-based systems. A cardiologist prescribing high-dose amiodarone for a refractory arrhythmia triggers a soft-limit alert in D1 that is routinely overridden, but the ML model recognises the clinical context (cardiac ICU, arrhythmia diagnosis, cardiology attending) and suppresses the non-actionable alert while escalating genuinely aberrant doses. For wrong-drug detection (Task E2), CDS-integrated alerts outperformed because drug allergy checking and interaction screening operate at a level of clinical specificity that dose-limit systems cannot reach. For cumulative dose tracking (Task E4), EHR auto-programming excelled because bidirectional communication enables the EHR to aggregate actual delivered volumes from pump feedback data, rather than relying on manual charting that typically underestimates cumulative exposure by 8-15%.

Table 3: Infusion Error Reduction Score (IERS) by Configuration (0-1)

Configuration	Interception	Specificity	Compliance	Workflow	Residual	IERS
D5: ML-Enhanced	0.940	0.920	0.920	0.900	0.940	0.922
D4: EHR Auto-Prog.	0.900	0.880	0.900	0.940	0.960	0.906
D3: CDS-Integrated	0.880	0.860	0.880	0.860	0.880	0.878
D2: Weight-Based	0.860	0.840	0.840	0.860	0.860	0.856
D1: Basic Limits	0.800	0.780	0.780	0.820	0.840	0.812

IERS weighted: interception 0.25, specificity 0.20, compliance 0.20, workflow 0.15, residual 0.20.

Table 4: Task-Specific Performance by Configuration

Configuration	E1: Wrong Dose	E2: Wrong Drug	E3: Wrong Rate	E4: Cumul. Dose
D1: Basic Limits	68.4%	12.1%	54.2%	41.8%
D2: Weight-Based	76.8%	14.6%	62.4%	48.6%
D3: CDS-Integrated	81.2%	72.4%	68.8%	56.2%

Configuration	E1: Wrong Dose	E2: Wrong Drug	E3: Wrong Rate	E4: Cumul. Dose
D4: EHR Auto-Prog.	88.6%	64.8%	86.2%	92.4%
D5: ML-Enhanced	96.8%	68.2%	88.4%	84.6%

Values represent interception rates (% of errors detected before reaching patient). E1-E4 measured independently.




Figure 1: Infusion Error Reduction Score (IERS) by Configuration




Figure 2: Interception Sensitivity vs. Alert Specificity




Figure 3: Residual Error Rate per 10,000 Infusions by Configuration




Figure 4: Multi-Dimensional Comparison -- Top 3 Configurations

5. Discussion

5.1 Solving Alert Fatigue Through Machine Intelligence

The ML-enhanced configuration's highest IERS confirms that the primary barrier to smart pump effectiveness is not technological capability but cognitive overload from poorly calibrated alerting. First-generation DERS treated all threshold violations equally, generating alerts for a paediatric intensivist prescribing milrinone at 1.2 mcg/kg/min (a routine dose that happens to exceed the conservative adult default limit) and for an inadvertent tenfold overdose of morphine with the same visual presentation. The ML model resolves this by learning contextual relevance: an alert for milrinone in a paediatric cardiac ICU with appropriate weight-based dosing has a significance score of 0.08 (suppressed), while an alert for the same drug at 10x the expected dose has a score of 0.97 (escalated with explanation). The 3.3-fold improvement in clinician compliance demonstrates that reducing noise restores attention to genuine safety signals -- a finding consistent with attention economics theory predicting that alert utility follows a concave function of alert volume.

5.2 The Complementary Value of EHR Integration

The EHR auto-programming configuration's lowest residual error rate highlights a fundamentally different error prevention mechanism: eliminating the opportunity for error rather than detecting errors after they occur. While DERS-based approaches (D1-D3, D5) intercept errors at the pump programming stage, auto-programming prevents them from entering the workflow at all by transmitting the verified order directly to the pump. The remaining errors under D4 arose not

from pump programming but from upstream pharmacy verification mistakes that propagated through the automated chain, suggesting that the next frontier in infusion safety lies in pharmacist order verification accuracy rather than pump-level intervention. The optimal deployment combines D4 (auto-programming to eliminate transcription errors) with D5 (ML-enhanced alerts to catch the remaining clinically significant dosing errors), achieving complementary protection across different error pathways.

5.3 Limitations

Three limitations affect the generalisability of these findings. First, the study was conducted within a single academic health system with mature EHR infrastructure and dedicated pharmacy informatics support -- community hospitals with fewer resources may achieve lower effectiveness for the more technologically demanding configurations (D4, D5). Second, the ML model was trained on data from the same institution, and performance may degrade when applied to institutions with different prescribing patterns, patient populations, and drug library configurations without local calibration. Third, the 36-month study period coincided with the COVID-19 pandemic, which introduced atypical medication usage patterns (high-dose anticoagulation, novel therapeutics) that may have influenced both error rates and alert performance in ways that do not reflect normal clinical operations.

6. Conclusion

6.1 Key Findings

ML-enhanced predictive alerts achieve the highest Infusion Error Reduction Score (0.922) by suppressing 64% of non-actionable alerts while maintaining 96.8% sensitivity for clinically significant errors, tripling clinician compliance rates through improved signal-to-noise ratio. EHR auto-programming achieves the lowest residual error rate (0.12 per 10,000 infusions) by eliminating manual transcription from the medication order-to-delivery chain. The optimal infusion safety architecture combines EHR auto-programming for error prevention with ML-enhanced alerting for error detection, addressing complementary pathways in the medication use process.

6.2 Future Directions

Three developments would advance smart infusion pump safety. First, federated learning across

multi-hospital networks would enable ML alert models trained on diverse clinical populations and prescribing patterns without sharing patient data, improving model generalisability beyond single-institution training sets. Second, closed-loop pharmacokinetic monitoring -- integrating real-time drug level sensors with pump rate adjustment algorithms -- would enable physiologically adaptive infusion dosing for drugs with narrow therapeutic indices such as vancomycin, aminoglycosides, and unfractionated heparin. Third, natural language processing applied to clinical notes and nursing documentation could detect early signs of adverse drug reactions that prompt proactive infusion rate adjustment before clinical deterioration. The SIPEAF evaluation framework, ML model code, drug library optimisation tools, and IERS calculator are available at sipeaf-safety.org under an open-source licence.

References

- Ancker, J. S. et al. (2017). Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system. *BMC Medical Informatics*, 17(1), 36.
- Bates, D. W. et al. (2020). The costs of adverse drug events in hospitalized patients. *JAMA*, 277(4), 307-311.
- Catlin, A. C. et al. (2015). Improving smart infusion pump drug libraries using DERS analytics. *Joint Commission Journal on Quality and Patient Safety*, 41(9), 396-403.
- Cho, I. et al. (2020). Effect of optimized alert thresholds on clinician responses. *Journal of the American Medical Informatics Association*, 27(6), 886-893.
- Dunford, B. B. et al. (2017). Smart pump implementation challenges: lessons from six hospitals. *American Journal of Health-System Pharmacy*, 74(11), 833-838.
- ECRI (2017). Top 10 Health Technology Hazards for 2017. ECRI Institute, Plymouth Meeting, PA.
- Husch, M. et al. (2005). Insights from the sharp end of intravenous medication errors. *Quality and Safety in Health Care*, 14(6), 426-431.
- ISMP (2019). 2018-2019 Targeted Medication Safety Best Practices for Hospitals. Institute for Safe Medication Practices.
- Keers, R. N. et al. (2013). Causes of medication administration errors in hospitals: a systematic review. *Drug Safety*, 36(11), 1045-1067.
- Lyons, I. et al. (2018). Infusion pump adverse events: the scale of the problem. *Drug Safety*, 41(12), 1203-1213.
- Melton, B. L. et al. (2019). Override rates for dose error reduction system alerts. *American Journal of Health-System Pharmacy*, 76(6), 391-398.
- Nuckols, T. K. et al. (2008). Cost-effectiveness of computerized intravenous medication safety system. *Annals of Internal Medicine*, 148(1), 45-56.
- Ohashi, K. et al. (2014). Benefits and risks of using smart pumps to reduce medication error rates: a systematic review. *Drug Safety*, 37(12), 1011-1020.
- Poon, E. G. et al. (2010). Effect of bar-code technology on the safety of medication administration. *New England Journal of Medicine*, 362(18), 1698-1707.
- Rothschild, J. M. et al. (2005). A controlled trial of smart infusion pumps to improve medication safety in critically ill patients. *Critical Care Medicine*, 33(3), 533-540.
- Taxis, K. and Barber, N. (2003). Ethnographic study of incidence and severity of intravenous drug errors. *BMJ*, 326(7391), 684.
- Westbrook, J. I. et al. (2011). Errors in the administration of intravenous medications and their association with information systems. *Journal of the American Medical Informatics Association*, 18(3), 373-379.
- Wilson, D. G. et al. (2004). Medication errors in paediatric practice: insights from an audit. *Archives of Disease in Childhood*, 89(12), 1128-1131.

Declarations

Funding

This research was supported by the Austrian Science Fund (FWF grant P34218-B) and the European Union Horizon 2020 programme (grant 965397, SafeInfusion). No funder influenced study design, data analysis, or the decision to publish.

Conflict of Interest

The authors declare no competing interests. No author has financial relationships with infusion pump manufacturers, EHR vendors, or medication safety consulting firms.

Data Availability Statement

The SIPEAF framework code, ML model training scripts, drug library analysis tools, and IERS calculator are available at sipeaf-safety.org under the MIT licence. Patient-level infusion event data cannot be shared due to institutional data governance policies.

Ethical Approval

This study analysed de-identified infusion pump log data and medication error reports under an institutional quality improvement designation. The Central European Tech University Ethics

Committee confirmed that formal research ethics approval was not required (waiver reference CETU-QI-2021-087).

Appendix A

Configuration Specifications and IERS Calculation Methodology

This appendix provides complete technical specifications for each DERS configuration, including drug library parameters, alert logic, EHR interface architecture, ML model details, and the IERS scoring rubric.

Part I -- Configuration Specifications

1a. D1: 468 drugs, 12 clinical profiles, fixed hard/soft limits per drug-profile pair

Industry-standard baseline

1b. D2: D1 + mandatory weight entry,
weight-normalised limits for 72% of ICU drugs
..... mg/kg dosing validation

2a. D3: D2 + renal/hepatic dose adjustment, allergy checking, interaction screening CDS rules from Lexicomp/Clinical Pharmacology

2b. D4: HL7 FHIR bidirectional interface,
auto-populated pump parameters from verified orders
..... Eliminates manual transcription

3a. D5: XGBoost (400 trees, 54 features), trained on 1.8M alert-outcome pairs Real-time significance scoring

3b. All: 6 hospitals, 36 months, 2.84M infusion events, independent pharmacist error adjudication Blinded evaluation

Part II -- IERS Scoring Rubric

4a. Interception: Fraction of clinically significant errors detected before delivery Adjudicated by clinical pharmacists

4b. Specificity: Fraction of alerts that are clinically actionable Dose changed appropriately upon alert

5a. Compliance: Fraction of appropriate alerts resulting in dose change Override = non-compliance

5b. Workflow: Inverse of alerts per 100 programming events Lower burden = higher score

6a. Residual Error: Confirmed errors per 10,000 infusions reaching patient Lower = better (inverse scoring)

6b. $IERS = 0.25 * Interception + 0.20 * Specificity + 0.20 * Compliance + 0.15 * Workflow + 0.20 * Residual$