International Journal of Innovative Technologies in Social Science

e-ISSN: 2544-9435

Operating Publisher
SciFormat Publishing Inc.
ISNI: 0000 0005 1449 8214

2734 17 Avenue SW,
Calgary, Alberta, T3E0A7,
Canada
+15878858911
editorial-office@sciformat.ca

ARTICLE TITLE	ARTIFICIAL INTELLIGENCE IN ANTIMICROBIAL STEWARDSHIP: CLINICAL APPLICATIONS, PREDICTIVE PERFORMANCE, AND IMPLEMENTATION CHALLENGES
DOI	https://doi.org/10.31435/ijitss.1(49).2026.5272
RECEIVED	07 January 2026
ACCEPTED	09 March 2026
PUBLISHED	20 March 2026
LICENSE	The article is licensed under a **Creative Commons Attribution 4.0 International License.**

ARTIFICIAL INTELLIGENCE IN ANTIMICROBIAL STEWARDSHIP: CLINICAL APPLICATIONS, PREDICTIVE PERFORMANCE, AND IMPLEMENTATION CHALLENGES

Oliwer Muller *(Corresponding Author, Email: oliwermuller@gmail.com)*
Kielce Hospital of St. Aleksandra Sp. z o.o., Kielce, Poland
ORCID ID: 0009-0005-6197-8461

Michał Babicz
Jan Kochanowski University, Kielce, Poland
ORCID ID: 0009-0007-7611-4695

Kamil Chudzicki
Wojewódzki Szpital Zespolony w Kielcach, Kielce, Poland
ORCID ID: 0009-0001-2734-7720

Wiktor Czyżewski
Independent Public Health Care Institution of the Ministry of the Interior and Administration in Kielce named after St. John Paul II, Kielce, Poland
ORCID ID: 0009-0003-0277-8203

Daria Danielczyk
Wojewódzki Szpital Zespolony w Kielcach, Kielce, Poland
ORCID ID: 0009-0002-4955-4883

Natalia Malatyńska
Jan Kochanowski University, Kielce, Poland
ORCID ID: 0009-0005-3824-8225

Anna Szot
Jan Kochanowski University, Kielce, Poland
ORCID ID: 0009-0003-2613-1068

Dominik Szydełko
Wojewódzki Szpital Zespolony w Kielcach, Kielce, Poland
ORCID ID: 0009-0002-9907-858X

Martyna Szymczyk
Jan Kochanowski University, Kielce, Poland
ORCID ID: 0009-0005-0772-7119

Paweł Żurek
Medical University of Warsaw, Warsaw, Poland
ORCID ID: 0009-0006-3023-5128

ABSTRACT

Artificial intelligence (AI) and machine learning (ML) are increasingly being proposed as tools to strengthen antimicrobial stewardship (AMS), particularly where antibiotic decisions depend on large volumes of clinical, microbiological, and prescribing data. This review synthesizes recent PubMed-indexed full-text evidence on AI in AMS, with emphasis on clinical application, predictive performance, implementation readiness, and governance. A structured narrative review was conducted using a targeted PubMed search updated on 8 March 2026. English-language peer-reviewed full-text papers published from 2017 through early 2026 were eligible when they addressed antimicrobial prescribing, susceptibility or resistance prediction, rapid diagnostic support, stewardship workflow, implementation, or ethical and organizational aspects of AI-enabled prescribing. Thirty-four papers were retained for structured synthesis, and two additional papers were used for contextual framing. The evidence indicates that AI is most mature in bounded stewardship tasks, especially individualized empiric prescribing, resistance prediction, rapid diagnostic interpretation, and prioritization of stewardship interventions. Newer open-access studies also show movement toward reusable stewardship datasets, real-time triage tools, and models for high-risk populations. However, the field remains dominated by retrospective, single-centre, and syndrome-specific studies, while external validation, workflow evaluation, and post-deployment governance are still limited. Current evidence therefore supports AI as a useful adjunct to stewardship teams rather than an autonomous replacement for clinical judgment.

KEYWORDS

Artificial Intelligence, Antimicrobial Stewardship, Antimicrobial Resistance, Machine Learning, Clinical Decision Support Systems, Antibiotic Prescribing

CITATION

Oliwer Muller, Michał Babicz, Kamil Chudzicki, Wiktor Czyżewski, Daria Danielczyk, Natalia Malatyńska, Anna Szot, Dominik Szydełko, Martyna Szymczyk, Paweł Żurek. (2026) Artificial Intelligence in Antimicrobial Stewardship: Clinical Applications, Predictive Performance, and Implementation Challenges. *International Journal of Innovative Technologies in Social Science*. 1(49). doi: 10.31435/ijitss.1(49).2026.5272

Introduction

Antimicrobial resistance (AMR) is no longer a distant microbiological concern. It now shapes everyday clinical practice by making common infections harder to treat, increasing the risk attached to surgery, oncology, transplantation, and intensive care, and raising the probability that delayed or inappropriate empirical therapy will translate into harm. Global burden estimates confirm that the scale of the problem is already substantial and that the human cost is likely to grow if antimicrobial use is not better governed in the coming decades (GBD 2021 Antimicrobial Resistance Collaborators, 2024).

In response to this pressure, antimicrobial stewardship (AMS) has become one of the central organizing principles of contemporary infectious disease care. Stewardship is broader than restriction. It is the disciplined effort to choose the right drug, in the right patient, at the right time, for the right duration, and then to revisit that decision as new information becomes available. This makes AMS inherently dynamic. It depends on microbiology, pharmacology, clinical context, institutional policy, and communication across teams. As a result, stewardship is both a technical and an organizational activity rather than a single prescribing rule (Dyar et al., 2017).

This is also why stewardship poses a difficult information problem. Antibiotic choices often need to be made before confirmatory culture results are available, under time pressure, and in patients with prior antibiotic exposure, incomplete records, or multiple competing risks. Even in data-rich settings, the relevant evidence may be distributed across electronic health records, historical microbiology, allergy documentation, ward-level resistance patterns, and evolving clinical observations. Earlier digital stewardship tools improved some process outcomes, but they often functioned primarily as static reminders, surveillance aids, or rule-based restriction systems. Their limitations created space for interest in AI-enabled decision support, which promises to use larger and more heterogeneous datasets in a more individualized way (Van Dort et al., 2022; Peiffer-Smadja et al., 2020).

Recent review literature suggests that AI in stewardship is moving from concept to application. Across infectious disease and stewardship-focused publications, AI and ML have been linked to individualized empiric prescribing, antimicrobial susceptibility prediction, earlier recognition of resistance, post-prescription review prioritization, rapid diagnostic interpretation, and workflow optimization. At the conceptual level, these systems are attractive because they offer a path away from purely population-level guidance toward patient-specific probability estimates that may better balance early adequacy of therapy against the long-term harms of broad-spectrum overuse (Chang & Chen, 2022; Theodosiou & Read, 2023; Rawson et al., 2024; Liu et al., 2024; Blechman & Wright, 2024; Cesaro et al., 2025).

Yet the field remains uneven. Technical studies often emphasize discrimination metrics, while implementation papers focus on usability, trust, and organizational fit. Ethical analyses raise concerns about fairness, liability, and the professional meaning of delegated decision support. Review-level syntheses increasingly show that the core question is no longer whether AI can produce attractive prediction results in selected settings, but whether those results can be translated into dependable stewardship value under real clinical conditions (Pinto-de-Sá et al., 2024; AlGain et al., 2025; Bosetti et al., 2025; Giacobbe et al., 2025).

That broader framing is especially relevant for a journal concerned with innovative technologies in social science. AI in AMS is not just a matter of algorithm design. It also concerns the organization of antibiotic decision-making, the division of responsibility between clinicians and software, the representativeness of local data, and the conditions under which hospitals and health systems actually adopt new tools. The present review was therefore designed to synthesize current PubMed-indexed full-text evidence on AI in AMS across four linked domains: clinical application, predictive value, implementation, and governance. It asks three practical questions: where AI appears most useful at present, what kind of evidence supports that view, and what barriers must be addressed before wider adoption is justified.

Methodology

Review design

This manuscript was prepared as a structured narrative review. That design was chosen deliberately. The relevant literature does not consist of one study genre or one stable outcome family that would support a single pooled quantitative approach. Instead, the field includes systematic reviews, meta-analyses, scoping reviews, conceptual papers, retrospective model-development studies, implementation surveys, qualitative ethics work, and early clinician-behavior experiments. A structured narrative design was therefore considered the most appropriate method for integrating these heterogeneous forms of evidence while still preserving transparency about how sources were identified, selected, and interpreted.

Search strategy

A targeted literature search was conducted in PubMed and updated on 8 March 2026. The search window covered publications from 1 January 2017 through 8 March 2026 so that the review would capture the period in which AI-enabled stewardship literature became clearly visible while still keeping the corpus bounded and reproducible. Search concepts were built around the main domains of the topic rather than around one overly restrictive Boolean string. The core search phrases included: "artificial intelligence in antimicrobial stewardship," "machine learning in antimicrobial stewardship," "clinical decision support in antimicrobial stewardship," "machine learning antibiotic resistance prediction," and "implementation of AI in antimicrobial stewardship." In practice, these seed searches were supplemented by manual screening of PubMed "similar articles," title-and-abstract review, backward citation checking from highly relevant papers, and a final targeted refresh for recent open full-text records from 2025-2026. To maintain continuity with the project's starting evidence pool, the PubMed-indexed papers already collated in the source file were also re-checked at full-text level, but only PubMed-indexed peer-reviewed papers were eligible for inclusion in the manuscript.

Table 1. Search framework and eligibility logic used in the review

Component	Operationalization in the present review
Database	PubMed
Search update	8 March 2026
Publication window	1 January 2017 to 8 March 2026
Main search domains	AI in AMS; ML in AMS; AI-CDSS in AMS; susceptibility/resistance prediction; implementation and ethics of AI-enabled prescribing
Screening steps	Plain-language seed search -> title/abstract screening -> full-text assessment -> PubMed "similar articles" check -> backward citation tracking -> final 2025-2026 open full-text refresh
Priority study types	Systematic reviews, meta-analyses, scoping reviews, original clinical or implementation studies, qualitative ethics work, dataset/infrastructure papers with direct AMS relevance, and high-value conceptual reviews
Inclusion criteria	English-language, peer-reviewed, full-text papers indexed in PubMed and directly relevant to antimicrobial prescribing, resistance prediction, rapid diagnostics, stewardship workflow, implementation, or governance
Exclusion criteria	Non-peer-reviewed material, conference abstracts without sufficient detail, editorials/news items without substantive analysis, duplicates, and AI-AMR papers without a clear stewardship or prescribing connection
Synthesis approach	Structured narrative synthesis; no new meta-analysis performed because study designs, endpoints, and tasks were too heterogeneous for meaningful pooling

As Table 1 indicates, the review logic was intentionally broad at the search stage but conservative at full-text inclusion, with preference given to clinically interpretable PubMed records that could inform real stewardship practice rather than AI-AMR literature in the abstract.

Eligibility criteria

Studies were eligible when they met four conditions. First, they had to be indexed in PubMed and available in full text. Second, they had to be published in English within the defined time window. Third, they had to address at least one stewardship-relevant decision domain: empiric prescribing, susceptibility or resistance prediction, rapid diagnostic support, post-prescription review, stewardship workflow, implementation, or ethical and organizational aspects of AI-supported antibiotic decision-making. Fourth, they had to provide either primary data or a substantial synthesis of the field.

Papers were excluded when their relevance to stewardship was only indirect. This applied particularly to studies on antimicrobial discovery, peptide design, basic molecular screening, or purely laboratory computation without a prescribing or stewardship endpoint. News coverage, brief opinion pieces, and technology commentary without analytic depth were also excluded. This conservative strategy was adopted because the broader AI-AMR literature is much larger than the stewardship-specific literature, and the aim of the present manuscript was to remain focused on clinically and organizationally meaningful antibiotic decision-making.

Study selection and prioritization

Screening was performed manually in two stages. First, titles and abstracts were reviewed for apparent relevance to stewardship-oriented antibiotic use. Second, potentially eligible records underwent full-text assessment. During full-text review, papers were prioritized when they offered one or more of the following: direct stewardship relevance, clinically interpretable outcomes, insight into implementation or ethics, recency, open-access availability, or high synthesis value. High-level evidence syntheses were retained even when some overlap existed in their underlying primary studies, because those papers contributed distinct analytic perspectives. After de-duplication, full-text review, and the final 2025-2026 refresh, 34 papers were retained for structured synthesis. Two additional papers were used for contextual framing of AMR burden and stewardship principles.

Data extraction and analytic framework

A structured evidence matrix was used to extract information from each included paper. The matrix captured publication year, setting, study type, data source, AI or digital method, stewardship task, main outcomes or performance measures, implementation relevance, and major limitations. Instead of organizing the review chronologically, the synthesis was arranged around six themes that reflected both the literature and the aims of the journal: the shape of the current evidence base; technical patterns and stewardship tasks; prescribing optimization and individualized decision support; resistance prediction, diagnostics, and stewardship prioritization; scaling stewardship through triage, reusable data assets, and special-population models; and implementation, ethics, and governance. This thematic structure helped reduce study-by-study repetition and made it possible to interpret technical findings within a broader clinical and organizational context.

Methodological appraisal

Because the corpus combined reviews, retrospective model studies, implementation research, and qualitative ethics work, no single formal appraisal tool could be applied meaningfully across the entire dataset. Methodological quality was therefore assessed narratively using a set of recurring questions: Was the target clinical or organizational problem clearly specified? Were the inputs and outcomes relevant to stewardship practice? Was model performance reported transparently and with some form of out-of-sample testing? Was external validation attempted or discussed? And did the paper address implementation, explainability, fairness, or workflow rather than treating prediction as an isolated end point? This form of appraisal does not replace formal risk-of-bias instruments, but it is suitable for a structured narrative review in a heterogeneous and fast-moving field.

Ethical considerations

The review used published literature only and did not involve new data collection from human participants or animals. Formal ethics approval was therefore not required.

Results

Shape of the current evidence base

The literature now contains enough material to support more than a purely speculative discussion of AI in stewardship, but the evidence remains uneven in depth, scope, and maturity. At the review level, the field has accelerated markedly in the last two years. Stewardship-specific syntheses now include a systematic review focused on AI in AMS, a broader systematic review on AI-driven stewardship and prescription optimization, a scoping review on antimicrobial prescribing and resistance, and quantitative syntheses examining predictive performance or comparison with traditional risk scoring systems. Together, these reviews suggest that the strongest evidence clusters around individualized empiric therapy, susceptibility or resistance prediction, rapid diagnostic support, and prioritization of stewardship attention, while implementation evidence remains thinner and more context-specific (Pinto-de-Sá et al., 2024; Harandi et al., 2025; Al Mazrouei et al., 2025; Pennisi et al., 2025; Pinto et al., 2025).

A second layer of the literature is more conceptual and translational. These papers do not always limit themselves to stewardship, but they are useful because they explain how infectious disease AI has evolved, what kinds of inputs these systems use, and why decision support in this area is both attractive and difficult. Across these publications, a consistent message emerges: the clinically plausible role of AI lies less in free-form automation and more in the structured integration of patient history, microbiology, prescribing patterns, and laboratory data around specific decision points. At the same time, these papers repeatedly warn that performance claims mean little if systems cannot be understood, monitored, and fitted into practice (Peiffer-Smadja et al., 2020; Theodosiou & Read, 2023; Liu et al., 2024; Blechman & Wright, 2024; Giacobbe et al., 2025; Cesaro et al., 2025).

The primary literature is narrower than the review literature and more skewed toward highly tractable tasks. Most original studies remain retrospective and single-centre. They frequently use urinary tract infection, structured susceptibility prediction, or well-defined post-prescription decisions as their proving ground. This concentration is understandable: these are areas where labels are clearer, data are more complete, and the stewardship question can be translated into a prediction task. However, it also means that the current literature may overrepresent situations in which algorithmic success is easier to achieve than in routine infectious disease practice. Studies addressing uncertain syndromes, cross-setting care, or multi-morbidity-driven prescribing complexity are still relatively uncommon (Hebert et al., 2020; Kanjilal et al., 2020; Bystritsky et al., 2020; Cai et al., 2023; Lin et al., 2024; Tran-The et al., 2024).

Table 2. Representative primary, implementation, and clinician-facing studies included in the synthesis

Study	Setting	Stewardship task	Main contribution
Hebert et al. (2020)	Hospitalized urinary tract infection	Individualized susceptibility prediction	Showed that patient-specific prediction can outperform static antibiogram logic when calibration is considered.
Kanjilal et al. (2020)	Outpatient uncomplicated urinary tract infection	Empiric antibiotic selection	Reduced second-line prescribing and lowered inappropriate therapy in a bounded outpatient use case.
Bystritsky et al. (2020)	Inpatients receiving broad-spectrum agents	Prediction of stewardship intervention need	Suggested that ML can help identify which patients are most likely to require stewardship review.
Cai et al. (2023)	Recurrent urinary tract infection	ANN-guided empiric treatment prediction	Reported strong sensitivity and specificity for predicting treatment efficacy in a recurrent-UTI context.
Cavallaro et al. (2023)	Resistance-aware prescribing	Explainable treatment matching	Demonstrated that interpretable models can reduce likely therapy mismatch while preserving transparency.
Bolton et al. (2024)	Hospital intravenous-to-oral switch decisions	Fair interpretable decision support	Showed that fairness-aware, clinically interpretable modelling can be built around a stewardship-relevant decision.
Lin et al. (2024)	MALDI-TOF-supported resistance detection	Rapid diagnostic support	Generated earlier resistance information than standard workflows in a clinically important resistant pathogen.
Huang et al. (2024)	Physicians using antibiotic AI-CDSS	Ethics and trust	Identified fairness, responsibility, and legal liability as key barriers to confident AI use in prescribing.
Tokgöz et al. (2024)	Hospital managers	Organizational implementation factors	Highlighted readiness, compatibility, training, and user-friendliness as practical adoption conditions.
Haredasht et al. (2025)	Inpatient urine, blood, and respiratory cultures	Personalized antibiograms	Demonstrated that EHR-linked models can support individualized resistance estimates across multiple specimen types.
Bolton et al. (2025)	Multicentre clinician-behavior experiment	AI-CDSS influence on prescribing	Showed that usability and selective behavioral influence matter more than theoretical acceptance alone.
Dutey-Magni et al. (2026)	Acute hospital postprescription review	Real-time prioritization of AMS attention	Demonstrated moderate predictive performance but meaningful improvement in identifying patients most likely to need therapy change.
Gallardo-Pizarro et al. (2026)	Febrile neutropenia in hematologic malignancy	Personalized empirical therapy in a high-risk population	Extended ML-supported AMS toward a vulnerable, high-stakes prescribing context.

As summarized in Table 2, the strongest original studies remain clustered around bounded prescribing tasks, but newer open-access work is clearly broadening the field toward stewardship triage, reusable data infrastructure, and higher-risk patient groups.

Resistance prediction, rapid diagnostics, and prioritization of stewardship attention

Prediction of susceptibility and resistance is the most technically mature AI application in AMS. Review-level evidence repeatedly identifies this domain as one of the clearest sources of signal. Quantitative syntheses report favorable pooled performance across conventional diagnostic and predictive metrics, although heterogeneity remains substantial and direct comparison across tasks is often difficult (Pennisi et al., 2025; Pinto et al., 2025). The reason this domain is attractive is straightforward: even modest improvements in anticipating resistance can support earlier narrowing, escalation, or targeting of therapy before full conventional reporting is complete.

Primary studies illustrate that point from different angles. Explainable resistance-aware models have shown that AI can estimate the probability of mismatch between a proposed treatment and likely susceptibility while also making the main drivers of risk visible to clinicians and stewardship teams (Cavallaro et al., 2023). Machine-learning approaches integrated with MALDI-TOF workflows suggest that some forms of resistance information can be generated meaningfully earlier than standard laboratory pipelines, which may shift the timing of stewardship action rather than simply reproducing the same information faster (Lin et al., 2024). Other studies show that AI can support the prioritization of stewardship effort itself, helping teams focus limited attention on prescriptions or patients where the probability of actionable intervention is higher (Bystritsky et al., 2020; Tran-The et al., 2024).

The importance of explainability is especially clear in this domain. Resistance prediction directly affects therapeutic adequacy, so stewardship teams need more than a black-box output. They require sufficient transparency to judge whether the model is relying on plausible drivers, whether the recommendation fits local epidemiology, and when human override is appropriate. In this respect, the resistance-prediction literature is useful not only because it reports encouraging performance, but also because it makes visible the conditions under which prediction becomes clinically usable rather than merely statistically impressive (Cavallaro et al., 2023; Bosetti et al., 2025; Cesaro et al., 2025).

Scaling stewardship through triage, reusable data assets, and special-population models

A particularly important recent shift is the move from one-off prediction tools toward systems that help stewardship teams allocate scarce specialist attention. Postprescription review is effective, but it is labor-intensive, and many hospitals do not have the staffing capacity to examine every antibiotic order with the same intensity. Dutey-Magni et al. (2026) addressed exactly this use case by training classifiers on real stewardship decisions and then prospectively validating them in real time. The best model achieved only moderate discrimination, but its clinical utility was still meaningful: if only the top 30% of antibiotic-treated inpatients could be reviewed, the model helped identify 68 of 145 patients needing therapy change or discontinuation, compared with 43 of 145 under random review. This kind of result is significant because it reframes success. A stewardship model does not need to be diagnostically perfect to improve the yield of finite pharmacist or infection-specialist time.

A second emerging line of work concerns reusable stewardship data infrastructure rather than isolated local prototypes. Haredasht et al. (2025) developed personalized antibiograms using 49,872 urine, blood, and respiratory infections derived from electronic health records. Their models achieved AUROCs of 0.74-0.78 across five antibiotics and highlighted prior resistance and prior antibiotic exposure as especially informative variables. The same research ecosystem also generated the Antibiotic Resistance Microbiology Dataset (ARMD), a de-identified EHR-linked microbiology resource specifically designed for AMR and stewardship research (Nateghi Haredasht et al., 2025). Although a dataset paper is not an intervention study, it matters for stewardship because limited benchmarking and poor comparability remain major barriers to external validation. Curated reusable datasets make it easier to test transportability, fairness, recalibration strategies, and cross-site robustness rather than relying only on bespoke local models.

A third development is the appearance of AI-supported prescribing work in higher-risk patient groups rather than only in common syndromes such as urinary tract infection. Gallardo-Pizarro et al. (2026) examined febrile neutropenia in patients with hematological malignancies and modeled empirical antibiotic selection using machine learning plus prior colonization with multidrug-resistant gram-negative bacilli. The importance of this paper lies not only in its proof-of-concept design, but in the clinical population it addresses. Stewardship in immunocompromised patients has a narrower margin for error than many outpatient or low-acuity use cases, so the ability of AI to contribute even theoretically in such settings suggests that the field is beginning to reach more clinically demanding contexts. At the same time, the study also illustrates why prospective confirmation remains essential before complex high-risk prescribing models can be generalized.

Implementation, clinician trust, and governance

The translational literature shows that implementation is not a secondary challenge to be solved after model development. It is part of the core evidence question. Earlier syntheses of digital stewardship interventions already indicated that computerized systems can reduce antibiotic consumption or improve appropriateness, but that their impact depends heavily on local context, system design, and workflow fit (Van Dort et al., 2022). The same principle appears even more strongly in recent AI-focused literature. Reviews now repeatedly identify usability, interoperability, training, governance, and trust as decisive conditions for real-world success (AlGain et al., 2025; Al Mazrouei et al., 2025; Tokgöz et al., 2023).

Managerial and clinician-facing studies make these constraints concrete. Survey evidence from hospital managers shows general openness to AI-enabled decision support, but also low self-perceived knowledge and a strong emphasis on implementation prerequisites such as technical compatibility, user-friendliness, and organizational support (Tokgöz et al., 2024). Qualitative work with physicians shows that doctors may welcome AI as support while still resisting the idea that responsibility for antibiotic choice can be delegated to a system whose training data, fairness, or legal status they do not fully understand (Huang et al., 2024). In a randomized multimethod study, clinicians exposed to AI-supported recommendations reported good usability but still exercised considerable discretion, indicating that acceptance of the tool did not translate into automatic compliance with its output (Bolton et al., 2025).

A further lesson from recent full-text literature is that implementation evidence is becoming more nuanced rather than uniformly optimistic. Earlier syntheses of digital stewardship interventions already suggested that computerized tools can improve prescribing quality, but that their effect depends strongly on how they are embedded within multidisciplinary stewardship processes rather than deployed as stand-alone technical products (Van Dort et al., 2022). Bosetti et al. (2025) make a similar point from the perspective of antimicrobial stewards themselves: computerized decision support systems often underperform not because the concept is flawed, but because alert burden, reactive design, poor timing, and weak integration into established care pathways undermine use. This is important for interpreting AI studies. When a model fails to change practice, the explanation may lie less in inadequate prediction than in the social and organizational architecture through which that prediction is delivered.

Recent review and perspective papers also show that implementation must be understood against a broader backdrop of health-data quality and informatics maturity. Blechman and Wright (2024) argue that electronic health record data are attractive for stewardship modelling precisely because they aggregate microbiology, prescribing, and patient-history variables, but those same data are frequently incomplete, inconsistently coded, and shaped by local workflow conventions. El Arab et al. (2025) extend that logic to infection prevention more broadly, emphasizing that AI rarely enters hospitals as an isolated intervention. It arrives as part of a larger digital ecosystem involving diagnostics, surveillance, infection prevention, and operational decision support. From a stewardship standpoint, this means that success depends partly on whether the hospital has enough informatics coherence to support data linkage, model updating, auditability, and user training over time.

The most recent ethical literature further suggests that not all forms of AI should be treated as equivalent in stewardship discourse. Rule-based or task-specific ML tools trained for explicit clinical questions are qualitatively different from open-ended generative systems. That distinction matters because current stewardship evidence is strongest for bounded prediction tasks, whereas large language models remain more vulnerable to inconsistency, prompt sensitivity, and opaque failure modes. Panda and Ghosh (2025) emphasize that transparency, privacy protection, fairness, and accountable human oversight remain core ethical requirements whenever AI contributes to infectious-disease diagnosis or therapy management. In practical terms, this means that the path from promising AI to acceptable stewardship infrastructure is likely to run through carefully governed, domain-specific tools rather than through broad conversational systems acting as quasi-autonomous prescribers.

These findings matter because they show that stewardship remains a human practice even when advanced analytics are introduced. Antibiotic decisions involve risk tolerance, institutional culture, professional judgment, and accountability to both the patient and the wider problem of resistance. A technically strong model can fail if it arrives too late, explains itself poorly, is not trusted, or disrupts the work patterns of already burdened teams. Conversely, a moderately strong model may still be useful if it is transparent, timely, and embedded well enough to support decisions clinicians already experience as difficult. In that sense, AI in AMS is best understood as a sociotechnical intervention rather than a software upgrade.

Table 3. Recurrent barriers to implementation and governance across the literature

Domain	Main issues reported in the literature	Why it matters for stewardship
Data quality and representativeness	Missing data, local dataset bias, non-transferability across institutions	Poor data foundations can make prescribing recommendations unreliable or unfair
External validation	Many models are developed in one site or one syndrome only	Good internal performance does not guarantee value in another hospital or patient population
Explainability	Black-box outputs are difficult to audit or trust in high-stakes prescribing	Stewardship teams need to understand when and why the system recommends a given action
Workflow fit	Alerts, timing, usability, and interoperability often determine uptake	Even accurate systems may fail if they arrive outside the real decision window
Professional accountability	Responsibility for following or overriding AI advice is often unclear	Ambiguity can reduce clinician trust and slow adoption
Training and organizational readiness	Users may be receptive but insufficiently prepared to work with AI-CDSS	Implementation depends on education, managerial support, and cultural acceptance
Fairness and governance	Non-local training data and opaque updating can reproduce inequity	Stewardship tools must be safe not only statistically but also socially and institutionally
Benchmarking and reusable data infrastructure	Publicly described cross-site datasets remain uncommon	Limited benchmarking makes external comparison, fairness testing, and recalibration harder

The cross-cutting issues condensed in Table 3 help explain why technically promising models so often struggle to reach durable clinical use.

Discussion

Main interpretation

The central conclusion of this review is that AI has established a credible but bounded place within AMS. The literature no longer supports dismissing the field as speculative. There are now multiple stewardship-focused reviews, quantitative syntheses, and primary studies demonstrating that AI can contribute to patient-specific empiric prescribing, resistance prediction, rapid diagnostic interpretation, and prioritization of stewardship effort. However, the same body of evidence also shows that usefulness is concentrated where the problem is tightly framed and the decision path is short. AI looks strongest when it narrows uncertainty around a clearly defined antibiotic question rather than when it attempts to absorb the full complexity of infectious disease care (Pinto-de-Sá et al., 2024; Harandi et al., 2025; Pennisi et al., 2025; Pinto et al., 2025).

From model performance to stewardship value

One recurring problem in the literature is the tendency to treat predictive strength as though it were equivalent to stewardship benefit. The two are related, but they are not the same. Stewardship teams need to know whether AI changes outcomes that matter to practice: fewer unnecessary broad-spectrum prescriptions, earlier optimization of therapy, lower mismatch risk, more efficient targeting of review, or more rapid transition from empiric to informed treatment. Some studies do address these questions directly, and those studies are especially valuable because they move beyond accuracy metrics into the terrain of real clinical use. Yet a large portion of the published literature still stops at internal validation or single-centre discrimination results. For that reason, strong performance statistics should be interpreted as necessary but incomplete evidence of stewardship value (Hebert et al., 2020; Kanjilal et al., 2020; Lin et al., 2024; Bolton et al., 2025).

This distinction also clarifies how current meta-analyses should be read. When pooled syntheses report favorable AUC, sensitivity, or predictive values, the conclusion is not that stewardship has solved its evidence problem. Rather, it is that a broad family of task-specific models has shown enough promise to justify more demanding forms of evaluation. The next evidentiary step is not endless repetition of retrospective proof-of-concept work. It is multicentre validation, calibration assessment, prospective testing, and demonstration that the model improves antibiotic decisions in settings where clinicians actually need help (Pennisi et al., 2025; Pinto et al., 2025; Bosetti et al., 2025).

Why implementation is a core scientific issue

The implementation literature makes one point with particular force: organizational fit is not an afterthought. Stewardship tools operate in environments shaped by handoff patterns, alert fatigue, documentation burden, managerial priorities, and professional hierarchies. A model that is elegant on paper may contribute little if it appears outside the real prescribing window, requires complicated data plumbing, or asks clinicians to trust an output that they cannot easily interpret. Conversely, a less ambitious model may still be useful if it is tightly aligned with existing workflow and supports decisions clinicians already find cognitively difficult (Van Dort et al., 2022; Tokgöz et al., 2023; Tokgöz et al., 2024).

This has implications for study design. If AI in AMS is investigated only as a prediction problem, the literature will continue to overstate technical progress and understate implementation risk. The inclusion of survey data, randomized multimethod studies, and qualitative ethics work is therefore not peripheral. It is necessary for understanding whether AI can function as a clinical technology rather than only as a research output. In this respect, stewardship may be one of the clearest medical domains in which to observe how algorithmic innovation depends on institutions, professions, and governance structures as much as on code.

Professional judgment, fairness, and responsibility

Another important finding is that the stewardship context exposes the limits of simplistic “AI versus clinician” framing. Doctors do not merely ask whether the system is accurate; they ask whether its recommendation is clinically plausible, whether its training data reflect their patients, whether they remain accountable after following it, and whether using it changes the moral balance of antibiotic prescribing. Those concerns are especially visible in work on explainability, fairness, and liability. Fair and interpretable modelling strategies have already been demonstrated in stewardship-relevant tasks, which is encouraging. At the same time, qualitative evidence shows that clinicians may remain reluctant to defer to AI in uncertain cases, particularly when narrower-spectrum recommendations feel risky or when legal responsibility is unclear (Bolton et al., 2024; Huang et al., 2024).

These issues become even more important when discussion turns to large language models. The available evidence suggests that current generative systems are less stable and less trustworthy for antibiotic guidance than tightly bounded ML tools trained for specific tasks. That does not mean generative systems have no place in infectious disease work, but it does mean that stewardship should treat them with caution. Antibiotic choices are high-stakes, time-sensitive, and ecologically consequential. In such a setting, fluent output is not a substitute for reliable task performance or accountable decision support (AlGain et al., 2025; Giacobbe et al., 2025).

Health-system implications and equity

The stewardship literature also points toward a broader health-system question. AMR is a global problem, but much of the AI evidence base has been produced in high-income institutions with mature digital infrastructure, strong microbiology services, and relatively rich data environments. This matters because transportability is not only statistical; it is infrastructural and social. Tools developed in one ecosystem may not perform similarly where laboratory workflows, prescribing norms, or data completeness differ. As digital-health reviews have argued, the contribution of AI to antimicrobial optimization cannot be separated from the capacity of institutions to maintain data quality, monitor performance, and support responsible use over time (Rawson et al., 2024).

That observation has an equity dimension. If AI-enabled stewardship tools are developed primarily in resource-rich settings and then promoted elsewhere without local validation, they may amplify rather than reduce existing inequalities in antimicrobial care. The burden of AMR is not distributed evenly, and the places with the highest burden are not always the places with the strongest digital infrastructure. Future research should therefore treat representativeness, local adaptation, and implementation feasibility as core components of stewardship evaluation rather than as optional later concerns (GBD 2021 Antimicrobial Resistance Collaborators, 2024; Harandi et al., 2025; Rawson et al., 2024).

A newer implication of the literature is that stewardship AI is beginning to require an infrastructure perspective rather than a model-by-model perspective. Papers on reusable microbiology datasets, large-scale stewardship analytics, and hospital-wide triage suggest that the next bottleneck may be less about inventing new algorithms and more about building stable data assets, shared benchmarks, and deployable evaluation pipelines (Tran-The et al., 2024; Nateghi Haredasht et al., 2025; Dutey-Magni et al., 2026). This matters because stewardship tools are unusually sensitive to local epidemiology. Without benchmarkable data and explicit recalibration strategies, even technically strong models may remain trapped within the institutions that produced them.

Recent broader reviews also reinforce that stewardship cannot be separated neatly from adjacent infection-prevention and ethical-governance systems. Integrative review work in hospital infection prevention shows that AI is increasingly embedded in broader digital infection-control ecosystems rather than isolated prescribing modules (El Arab et al., 2025). At the same time, more explicitly normative work now argues that the ethical use of AI in infectious-disease diagnosis and therapeutic stewardship depends on maintaining transparency, privacy protection, accountability, and meaningful human oversight (Panda & Ghosh, 2025). These papers do not replace stewardship-specific evidence, but they do widen the frame of interpretation: the question is no longer only whether a model predicts well, but whether the institution can use it responsibly, monitor it over time, and justify its recommendations to clinicians and patients.

Another reason to avoid overreading current performance results is that stewardship decisions are highly path-dependent. A recommendation is only useful if it appears at a time when the clinician can still act on it, if the local formulary or diagnostic pathway can support the alternative, and if the recommendation is legible enough to justify deviating from usual practice. This helps explain why some of the strongest early AI successes have emerged in tightly bounded questions such as urinary tract infection therapy, antibiogram personalization, or intravenous-to-oral switch review. Those tasks offer clear decision windows, measurable outcomes, and comparatively structured data. More diffuse syndromes, by contrast, may remain difficult not because AI is irrelevant, but because the surrounding care pathway is less standardized and therefore harder to augment reliably.

This review also suggests that future publication quality in the field will increasingly depend on methodological transparency beyond the algorithm itself. Readers now need to know not only which model performed best, but how missing data were handled, what the temporal validation design looked like, whether recalibration was considered, how end users were involved, and what happened when the model was tested prospectively or behaviorally. Several of the newest full-text studies are valuable precisely because they begin to answer these questions rather than only reporting another internally validated classifier. If that trend continues, the stewardship literature may shift from proof-of-concept enthusiasm toward a more mature science of deployment.

Strengths, limitations, and research priorities

This review has several strengths. It focuses on full-text, PubMed-indexed, peer-reviewed literature; it integrates primary studies with higher-level syntheses; and it treats implementation and governance as part of the evidence base rather than as background commentary. It also has limitations. The review is narrative rather than systematic in the PRISMA sense, even though its search and selection process was structured. Only one primary database was used, although citation tracking broadened coverage. The field itself is heterogeneous, and the evidence window is moving quickly, which means that any synthesis of AI in stewardship will require updating as new prospective studies appear.

Despite those limitations, the direction of future work is relatively clear. Multicentre validation should become routine. Outcomes should be reported in ways that connect prediction more directly with stewardship practice, such as time to optimal therapy, de-escalation success, unnecessary broad-spectrum exposure, or efficient allocation of review effort. More implementation science is needed, including studies of trust, explanation design, cost, training, and post-deployment monitoring. Finally, governance frameworks should be developed in parallel with model development rather than afterwards. The field has reached the point at which better algorithms alone are unlikely to solve the most important translational problems.

Conclusions

AI has become a serious, evidence-backed topic within antimicrobial stewardship. Current literature supports its use as an adjunctive tool in narrowly defined stewardship tasks, especially individualized empiric prescribing, susceptibility or resistance prediction, rapid diagnostic support, and prioritization of review. In those areas, AI can make existing clinical and microbiological data more actionable and may help reduce some of the uncertainty that drives unnecessary or poorly targeted antibiotic exposure.

At the same time, the evidence base remains more convincing technically than operationally. Many studies are retrospective, single-centre, and syndrome-specific; implementation work is growing but still modest; and ethical as well as governance questions remain unresolved. For that reason, the most defensible conclusion is not that AI is ready to replace stewardship teams, but that it can strengthen them when models are externally validated, appropriately explained, and embedded in real workflows with human oversight. The next phase of progress will depend less on headline algorithmic novelty and more on careful translation into the social, organizational, and clinical reality of antibiotic decision-making.

Conflicts of Interest: No conflicts of interest to declare.

REFERENCES

1. Al Mazrouei, N., Elnour, A. A., Badi, S., Alsulami, F. T., Saeed, A. A. M., Al-Kubaisi, K. A., Menon, V., Khidir, I. Y., Ismail, M., Mahagoub, M. M. O., Ramisetti, A. K., & Elkarib, R. (2025). The impact of artificial intelligence on the prescribing, selection, resistance, and stewardship of antimicrobials: A scoping review. *BMC Infectious Diseases, 26*, 222. https://doi.org/10.1186/s12879-025-12320-4
2. AlGain, S., Marra, A. R., Kobayashi, T., Marra, P. S., Celeghini, P. D., Hsieh, M. K., Shatari, M. A., Althagafi, S., Alayed, M., Ranavaya, J. I., Boodhoo, N. A., Meade, N. O., Fu, D., Sampson, M. M., Rodriguez-Nava, G., Zimmet, A. N., Ha, D., Alsuhaibani, M., Huddleston, B. S., & Salinas, J. L. (2025). Can we rely on artificial intelligence to guide antimicrobial therapy? A systematic literature review. *Antimicrobial Stewardship & Healthcare Epidemiology, 5*(1), e90. https://doi.org/10.1017/ash.2025.47
3. Blechman, S. E., & Wright, E. S. (2024). Applications of machine learning on electronic health record data to combat antibiotic resistance. *The Journal of Infectious Diseases, 230*(5), 1073–1082. https://doi.org/10.1093/infdis/jiae348
4. Bolton, W. J., Wilson, R., Gilchrist, M., Georgiou, P., Holmes, A., & Rawson, T. M. (2024). Personalising intravenous to oral antibiotic switch decision making through fair interpretable machine learning. *Nature Communications, 15*(1), 506. https://doi.org/10.1038/s41467-024-44740-2
5. Bolton, W. J., Wilson, R., Gilchrist, M., Georgiou, P., Holmes, A., & Rawson, T. M. (2025). The impact of artificial intelligence-driven decision support on uncertain antimicrobial prescribing: A randomised, multimethod study. *The Lancet Digital Health, 7*(11), 100912. https://doi.org/10.1016/j.landig.2025.100912
6. Bosetti, D., Grant, R., & Catho, G. (2025). Computerized decision support for antimicrobial prescribing: What every antibiotic steward should know. *Antimicrobial Stewardship & Healthcare Epidemiology, 5*(1), e210. https://doi.org/10.1017/ash.2025.10091
7. Bystritsky, R. J., Beltran, A., Young, A. T., Wong, A., Hu, X., & Doernberg, S. B. (2020). Machine learning for the prediction of antimicrobial stewardship intervention in hospitalized patients receiving broad-spectrum agents. *Infection Control & Hospital Epidemiology, 41*(9), 1022–1027. https://doi.org/10.1017/ice.2020.213
8. Cai, T., Anceschi, U., Prata, F., Collini, L., Brugnolli, A., Migno, S., Rizzo, M., Liguori, G., Gallelli, L., Wagenlehner, F. M. E., Bjerklund Johansen, T. E., Montanari, L., Palmieri, A., & Tascini, C. (2023). Artificial intelligence can guide antibiotic choice in recurrent UTIs and become an important aid to improve antimicrobial stewardship. *Antibiotics, 12*(2), 375. https://doi.org/10.3390/antibiotics12020375
9. Cavallaro, M., Moran, E., Collyer, B., McCarthy, N. D., Green, C., & Keeling, M. J. (2023). Informing antimicrobial stewardship with explainable AI. *PLOS Digital Health, 2*(1), e0000162. https://doi.org/10.1371/journal.pdig.0000162
10. Cesaro, A., Hoffman, S. C., Das, P., & de la Fuente-Nunez, C. (2025). Challenges and applications of artificial intelligence in infectious diseases and antimicrobial resistance. *npj Antimicrobials and Resistance, 3*(1), 2. https://doi.org/10.1038/s44259-024-00068-x
11. Chang, A., & Chen, J. H. (2022). BSAC Vanguard Series: Artificial intelligence and antibiotic stewardship. *Journal of Antimicrobial Chemotherapy, 77*(5), 1216–1217. https://doi.org/10.1093/jac/dkac096
12. Dutey-Magni, P. F., Brown, M., Harris, S., Curtis, C., Dobson, R., Chowdhury, H., Cawthorn, A., De, S., Stone, N., Cooper, J., & Shallcross, L. (2026). Supervised machine learning to identify hospital inpatients needing a change of antibiotic therapy in real time: Preclinical diagnostic evaluation and feasibility study. *Open Forum Infectious Diseases, 13*(1), ofaf721. https://doi.org/10.1093/ofid/ofaf721

13. Dyar, O. J., Huttner, B., Schouten, J., Pulcini, C., & ESGAP (ESCMID Study Group for Antimicrobial Stewardship). (2017). What is antimicrobial stewardship? *Clinical Microbiology and Infection, 23*(11), 793–798. https://doi.org/10.1016/j.cmi.2017.08.026
14. El Arab, R. A., Almoosa, Z., Alkhunaizi, M., Abuadas, F. H., & Somerville, J. (2025). Artificial intelligence in hospital infection prevention: An integrative review. *Frontiers in Public Health, 13*, 1547450. https://doi.org/10.3389/fpubh.2025.1547450
15. Gallardo-Pizarro, A., Teijón-Lumbreras, C., Aiello, T. F., Monzó-Gallo, P., Martinez-Urrea, A., Terrones-Campos, C., Cuervo, G., Morata, L., Castro, P., Casals, C., Espasa, M., Esteve, J., Nicolas, J. M., Mensa, J., Soriano, A., & Garcia-Vidal, C. (2026). Towards personalised empirical antibiotic therapy in febrile neutropenia: A theoretical model based on machine learning and prior colonisation with multidrug-resistant gram-negative bacilli—A retrospective proof-of-concept cohort study. *Therapeutic Advances in Infectious Disease, 13*, 20499361251411061. https://doi.org/10.1177/20499361251411061
16. GBD 2021 Antimicrobial Resistance Collaborators. (2024). Global burden of bacterial antimicrobial resistance 1990–2021: A systematic analysis with forecasts to 2050. *The Lancet, 404*(10459), 1199–1226. https://doi.org/10.1016/S0140-6736(24)01867-1
17. Giacobbe, D. R., Guastavino, S., Marelli, C., Murgia, Y., Mora, S., Signori, A., Rosso, N., Giacomini, M., Campi, C., Piana, M., & Bassetti, M. (2025). Antibiotics and artificial intelligence: Clinical considerations on a rapidly evolving landscape. *Infectious Diseases and Therapy, 14*(3), 493–500. https://doi.org/10.1007/s40121-025-01114-5
18. Harandi, H., Shafaati, M., Salehi, M., Roozbahani, M. M., Mohammadi, K., Akbarpour, S., Rahimnia, R., Hassanpour, G., Rahmani, Y., & Seifi, A. (2025). Artificial intelligence-driven approaches in antibiotic stewardship programs and optimizing prescription practices: A systematic review. *Artificial Intelligence in Medicine, 162*, 103089. https://doi.org/10.1016/j.artmed.2025.103089
19. Haredasht, F. N., Maddali, M. V., Ma, S. P., Chang, A., Kim, G. Y. E., Banaei, N., Deresinski, S., Goldstein, M. K., Asch, S. M., & Chen, J. H. (2025). Enhancing antibiotic stewardship: A machine learning approach to predicting antibiotic resistance in inpatient care. *AMIA Annual Symposium Proceedings, 2024*, 857–864.
20. Hebert, C., Gao, Y., Rahman, P., Dewart, C., Lustberg, M., Pancholi, P., Stevenson, K., Shah, N. S., & Hade, E. M. (2020). Prediction of antibiotic susceptibility for urinary tract infection in a hospital setting. *Antimicrobial Agents and Chemotherapy, 64*(7), e02236-19. https://doi.org/10.1128/AAC.02236-19
21. Huang, Z., Lim, H. Y.-F., Ow, J. T., Sun, S. H.-L., & Chow, A. (2024). Doctors' perception on the ethical use of AI-enabled clinical decision support systems for antibiotic prescribing recommendations in Singapore. *Frontiers in Public Health, 12*, 1420032. https://doi.org/10.3389/fpubh.2024.1420032
22. Kanjilal, S., Oberst, M., Boominathan, S., Zhou, H., Hooper, D. C., & Sontag, D. (2020). A decision algorithm to promote outpatient antimicrobial stewardship for uncomplicated urinary tract infection. *Science Translational Medicine, 12*(568), eaay5067. https://doi.org/10.1126/scitranslmed.aay5067
23. Lin, T.-H., Chung, H.-Y., Jian, M.-J., Chang, C.-K., Lin, H.-H., Yu, C.-M., Perng, C.-L., Chang, F.-Y., Chen, C.-W., Chiu, C.-H., & Shang, H.-S. (2024). Artificial intelligence-clinical decision support system for enhanced infectious disease management: Accelerating ceftazidime-avibactam resistance detection in *Klebsiella pneumoniae*. *Journal of Infection and Public Health, 17*(10), 102541. https://doi.org/10.1016/j.jiph.2024.102541
24. Liu, G.-Y., Yu, D., Fan, M.-M., Zhang, X., Jin, Z.-Y., Tang, C., & Liu, X.-F. (2024). Antimicrobial resistance crisis: Could artificial intelligence be the solution? *Military Medical Research, 11*(1), 7. https://doi.org/10.1186/s40779-024-00510-1
25. Nateghi Haredasht, F., Amrollahi, F., Maddali, M. V., Marshall, N., Ma, S. P., Cooper, L. N., Johnson, A. O., Wei, Z., Medford, R. J., Kanjilal, S., Banaei, N., Deresinski, S., Goldstein, M. K., Asch, S. M., Chang, A., & Chen, J. H. (2025). Antibiotic resistance microbiology dataset (ARMD): A resource for antimicrobial resistance from EHRs. *Scientific Data, 12*(1), 1299. https://doi.org/10.1038/s41597-025-05649-7
26. Panda, P. K., & Ghosh, S. (2025). Ethical use of AI in infectious diagnostic decision and therapeutic stewardship. *IDCases, 42*, e02356. https://doi.org/10.1016/j.idcr.2025.e02356
27. Peiffer-Smadja, N., Rawson, T. M., Ahmad, R., Buchard, A., Georgiou, P., Lescure, F.-X., Birgand, G., & Holmes, A. H. (2020). Machine learning for clinical decision support in infectious diseases: A narrative review of current applications. *Clinical Microbiology and Infection, 26*(5), 584–595. https://doi.org/10.1016/j.cmi.2019.09.009
28. Pennisi, F., Pinto, A., Ricciardi, G. E., Signorelli, C., & Gianfredi, V. (2025). Artificial intelligence in antimicrobial stewardship: A systematic review and meta-analysis of predictive performance and diagnostic accuracy. *European Journal of Clinical Microbiology & Infectious Diseases, 44*(3), 463–513. https://doi.org/10.1007/s10096-024-05027-y
29. Pinto, A., Pennisi, F., Ricciardi, G. E., Signorelli, C., & Gianfredi, V. (2025). Evaluating the impact of artificial intelligence in antimicrobial stewardship: A comparative meta-analysis with traditional risk scoring systems. *Infectious Diseases Now, 55*(5), 105090. https://doi.org/10.1016/j.idnow.2025.105090
30. Pinto-de-Sá, R., Sousa-Pinto, B., & Costa-de-Oliveira, S. (2024). Brave new world of artificial intelligence: Its use in antimicrobial stewardship—A systematic review. *Antibiotics, 13*(4), 307. https://doi.org/10.3390/antibiotics13040307

31. Rawson, T. M., Zhu, N., Galiwango, R., Cocker, D., Islam, M. S., Myall, A., Vasikasin, V., Wilson, R., Shafiq, N., Das, S., & Holmes, A. H. (2024). Using digital health technologies to optimise antimicrobial use globally. *The Lancet Digital Health, 6*(12), e914–e925. https://doi.org/10.1016/S2589-7500(24)00198-5
32. Theodosiou, A. A., & Read, R. C. (2023). Artificial intelligence, machine learning and deep learning: Potential resources for the infection clinician. *Journal of Infection, 87*(4), 287–294. https://doi.org/10.1016/j.jinf.2023.07.006
33. Tokgöz, P., Hafner, J., & Dockweiler, C. (2023). Factors influencing the implementation of decision support systems for antibiotic prescription in hospitals: A systematic review. *BMC Medical Informatics and Decision Making, 23*(1), 27. https://doi.org/10.1186/s12911-023-02124-4
34. Tokgöz, P., Krayter, S., Hafner, J., & Dockweiler, C. (2024). Decision support systems for antibiotic prescription in hospitals: A survey with hospital managers on factors for implementation. *BMC Medical Informatics and Decision Making, 24*(1), 96. https://doi.org/10.1186/s12911-024-02490-7
35. Tran-The, T., Heo, E., Lim, S., Suh, Y., Heo, K.-N., Lee, E. E., Lee, H.-Y., Kim, E. S., Lee, J.-Y., & Jung, S. Y. (2024). Development of machine learning algorithms for scaling-up antibiotic stewardship. *International Journal of Medical Informatics, 181*, 105300. https://doi.org/10.1016/j.ijmedinf.2023.105300
36. Van Dort, B. A., Penm, J., Ritchie, A., & Baysari, M. T. (2022). The impact of digital interventions on antimicrobial stewardship in hospitals: A qualitative synthesis of systematic reviews. *Journal of Antimicrobial Chemotherapy, 77*(7), 1828–1837. https://doi.org/10.1093/jac/dkac112