Beyond Forward Chaining: Logical Perspectives on Clinical Infection Monitoring of BSIs

Julia LIEPOLD[a,b,1], Leonhard HAUPTFELD[b], Moritz GROB[b,c], Andrea RAPPELSBERGER[c] and Klaus-Peter ADLASSNIG[b,c]

[a]*Institute for Logic and Computation, TU Wien, 1040 Vienna, Austria*

[b]*Medexter Healthcare, Borschkegasse 7/5, 1090 Vienna, Austria*

[c]*Medical University of Vienna, Center for Medical Data Science, Institute of Artificial Intelligence, Spitalgasse 23, 1090 Vienna, Austria*

ORCiD ID: Julia Liepold 0009-0004-5416-6417, Leonhard Hauptfeld 0009-0006-5902-4666, Moritz Grob 0009-0006-7867-9465, Andrea Rappelsberger 0009-0001-0345-4011, Klaus-Peter Adlassnig 0009-0008-5046-1549

Abstract. Healthcare-associated infections (HAIs), particularly bloodstream infections (BSIs), require timely and reliable detection. Typically, rule-based clinical decision support systems use forward chaining to propagate patient data through knowledge-based systems. While effective, this data-driven approach can become computationally costly in data-rich settings and obscures the underlying reasoning structure. We analyze infection monitoring from a computational perspective by formalizing rule-based detection as a fixed-point computation and examining its complexity. As an alternative, we consider goal-driven reasoning, which focuses inference on specific diagnostic hypotheses, and formulate BSI detection as a satisfiability problem in propositional logic. This perspective supports more efficient, interpretable, and uncertainty-aware approaches to clinical decision support.

Keywords. clinical reasoning, bloodstream infections, forward & backward chaining, satisfiability problem (SAT)

1. Introduction

Clinical diagnosis involves not only data interpretations, but also structured reasoning. Foundational work in biomedical computing and artificial intelligence in medicine conceptualizes clinical decision-making as a synthesis of information theory, Bayesian reasoning, decision analysis, and Boolean logic [1]. This highlights that diagnostic systems are not neutral tools, but explicit realizations of particular reasoning paradigms. Modern research continues to emphasize that clinical reasoning is inherently multifaceted, combining intuitive pattern recognition with more structured, hypothesis-driven approaches [2].

Consequently, the choice of computational model directly shapes evidence interpretation and clinical decision support systems, enabling large-scale diagnostic

[1] Corresponding Author: Julia Liepold; E-mail: jl@medexter.com

reasoning. Systems such as MONI (Monitoring of Nosocomial Infections), formerly deployed at University Hospital Vienna, implement rule-based inference using Medical Logic Modules encoded in Arden Syntax, an HL7 International standard for medical knowledge representation and processing [3]. These modules act as building blocks, enabling the construction of detection pipelines for healthcare-associated infections (HAIs). Among them, bloodstream infections (BSIs) are particularly severe, as pathogens entering the bloodstream can lead to systemic inflammation, septic shock, and high mortality.

The aim of this study is to formalize diagnostic reasoning for BSIs using a rule-based framework, building on prior pipeline approaches successfully applied to pneumonia detection [4], and to explore extensions to fuzzy and weighted reasoning for efficient inference.

2. Methods

Extending the framework introduced in [4], the detection of BSIs was implemented using the same pipeline method. This approach was chosen to facilitate comparability. According to the European Centre for Disease Prevention and Control the definition of BSI is:

- at least one positive blood culture for a recognized pathogen (P)
  OR
- both of the following:
  - at least one of the following symptoms: fever (F), chills (C), hypotension (H)
  - at least two positive blood cultures for a common skin contaminant (S)

Skin contaminants are further defined as one of the following: *coagulase-negative staphylococci*, *Micrococcus spp., Propionibacterium acnes, Bacillus spp., Corynebacterium spp.*, and are taken from two separate blood samples, usually within 48 hours [5].

From a computational perspective, this definition can be interpreted as a set of logical conditions over clinical observations, making it a suitable candidate for rule-based inference. Two main types of inference engines are commonly distinguished [6]: forward- chaining, a data- driven approach, and backward- chaining, a goal- driven approach. These paradigms are foundational in rule- based reasoning and expert systems [6] and form the basis for analytical comparisons in the following sections.

2.1. Data-driven approach

Data-driven reasoning, which is often achieved through forward chaining, serves as the basis of many rule-based clinical decision support systems, likewise in Medexter's MONI pipeline system [4]. In this process, inference proceeds from available observations toward higher-level conclusions, starting from known facts such as patient data. Rules are evaluated iteratively to derive new information until no further conclusions can be drawn. To show this, we used a pipeline editor to conceptualize the detection workflow for BSIs [7], as depicted in Figure 1.

The algorithm follows a recursive pattern: when a node is evaluated, its results are forwarded to downstream nodes, potentially enabling their execution once all required inputs are available. Auxiliary nodes are used for static input values and are resolved on demand, ensuring that dependent nodes receive all necessary information before being triggered.

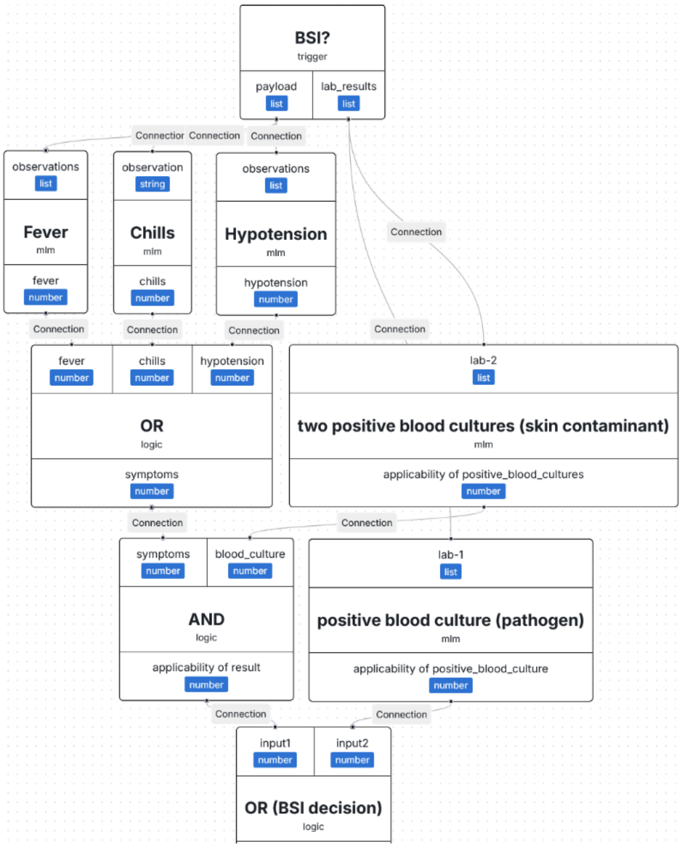

Figure 1. Pipeline-based implementation of fuzzy bloodstream infection (BSI) detection using forward-chaining.

2.2. Goal-driven approach

In contrast, goal-driven reasoning is usually realized via backward chaining, which inverts the direction of reasoning by beginning from a hypothesis and attempting to justify it by using available data [1,6]. Early expert systems, such as MYCIN [8], exemplify this model of reasoning where diagnostic conclusions are treated as goals to be proven rather than outcomes to be derived. Instead of attempting to propagate all accessible facts through the rule set, backward chaining selectively evaluates only those rules that are relevant to the query at hand. This effectively creates a search through the space of possible justifications, where rules are recursively expanded into subgoals until they can be matched against known facts [1].

3. Results

Having explained those concepts, we now examine their computational behavior and structural implications. Forward chaining is an iterative fixed-point computation over the rule set. This approach is sound and complete with respect to the underlying rule system [6]. Let R denote the number of rules and F the number of currently known facts. In the worst case, each newly derived fact may trigger the re-evaluation of all rules, resulting in a per-iteration complexity of $O(R * F)$. Since forward chaining incrementally derives new facts, potentially one at a time, up to $O(F)$ iterations can be required until convergence. In a densely connected clinical rule system, where intermediate results repeatedly propagate through the network, this leads to an overall complexity approaching $O(R * F^2)$ in practice.

In contrast, backward chaining evaluates only those rules relevant to a specific diagnostic goal. Its complexity is determined by the depth and branching of the proof tree rather than the overall size of the knowledge base. If we now define R as the number of rules relevant to a specific goal (i.e., the branching factor) and d as the depth of the proof tree, then the complexity of backward chaining amounts to $O(R^d)$. While theoretically exponential, d is typically small in clinical scenarios, making goal-driven reasoning more efficient when only a limited set of hypotheses (such as "does the patient have a BSI?") must be evaluated. Techniques such as memoization, pruning, and heuristic rule ordering may further reduce redundant computations [1,2].

This goal-directed view aligns naturally with a structural representation of diagnostic knowledge as an AND–OR tree, where conjunctions and disjunctions correspond to AND and OR nodes, respectively [6]. Methods, such as proof-number search [9], are available to guide the search and reduce the number of paths to be considered. Such graphs then can be directly translated into propositional logic, making it possible to reformulate diagnostic inference as a satisfiability problem (SAT). A diagnosis holds if the corresponding logical formula is satisfied by the observed data. If we extend this approach to fuzzy logic, we replace the Boolean operators with their fuzzy counterparts (e.g., t-norms and t-conorms), yielding a soft constraint system. In this view, diagnostic inference is equivalent to solving a fuzzy satisfiability or weighted Max-SAT problem, where each constraint carries an expert-defined weight reflecting its applicability, and the goal is to maximize the weighted sum of satisfied constraints. In our BSI example, using the assigned letters from the definition in Section 2, we get the formula $P \vee \big((F \vee C \vee H) \wedge S\big)$. Converting this to conjunctive normal form (the typical SAT input format) leaves us with two separate clauses:

1. $P \vee F \vee C \vee H$ and
2. $P \vee S$

Splitting the formula into sub-formulas is useful because it allows SAT solvers to track dependencies and apply optimizations, like clause learning and conflict-driven backtracking, more efficiently. Related extensions can also be framed in terms of probabilistic logic programming frameworks, such as ProbLog. While SAT is NP-complete, modern and highly optimized solver techniques, like heuristic search or sophisticated pruning, are directly applicable. This enables focusing the search directly on relevant diagnostic paths, providing an efficient and structured alternative to iterative fixed-point evaluation [6,10].

4. Discussion

It is important to note that the above reformulations are theoretical in nature and aim to uncover the underlying logical structure of infection monitoring rather than to propose a fully implemented alternative. The example considered in this work is a deliberately simplified version to provide conceptual understanding. Nevertheless, it is worth noting that this perspective highlights a valuable opportunity: clinical decision support systems might directly benefit from advancements in formal logic and satisfiability solving. Extending systems to fuzzy logic enables the natural representation of uncertainties in clinical data while retaining computational advantages [2,10].

Looking forward, SAT-based approaches could further enhance reasoning by identifying pathognomonic patterns (key symptom combinations or rule sets that directly infer a diagnosis), enabling faster conclusions with interpretable explanations.

5. Conclusion

The analysis of the detection of BSIs illustrates the fundamental trade-offs between data-driven and goal-driven reasoning in clinical decision support systems. Representing diagnostic rules as AND–OR graphs exposes their logical structure and enables reformulation as a propositional satisfiability problem. By viewing diagnostic problems as a logical framework, state-of-the-art SAT solver methods become available and offer new ways to reason about clinical data that are both computationally efficient and structurally transparent.

Overall, such interdisciplinary, logic-based methods combine the efficiency of goal-directed search with the flexibility to handle graded or uncertain evidence, offering a promising framework for next-generation infection monitoring systems.

References

[1] Ledley RS, Lusted LB. Reasoning Foundations of Medical Diagnosis. Science. 1959;130(3366):9-21. doi: 10.1126/science.130.3366.9

[2] Greengrass CJ. Transforming clinical reasoning—the role of AI in supporting human cognitive limitations. Frontiers in Digit Health. 2026;7:1715440. doi: 10.3389/fdgth.2025.1715440

[3] Health Level Seven International. Health Level Seven Arden Syntax for Medical Logic Systems, Version 3.0 STU2 [Internet]. Ann Arbor (MI), USA: HL7 International; 2023 Available from: https://www.hl7.org/implement/standards/product_brief.cfm?product_id=639 [cited 2026 March 20]

[4] Hauptfeld L, Rappelsberger A, Adlassnig KP. Infection Control Through Clinical Pipelines Built with Arden Syntax MLM Building Blocks. Studies in health technology and informatics. 2024;313:167-72. doi: 10.3233/SHTI240032

[5] European Centre for Disease Prevention and Control. Protocol for the surveillance of healthcare-associated infections and prevention indicators in European intensive care units. Stockholm, Sweden: ECDC; 2025. Available from: https://www.ecdc.europa.eu/sites/default/files/documents/protocol-surveillance-healthcare-associated-infections-intensive-care.pdf [cited 2026 March 20]

[6] Russell S, Norvig P. Artificial Intelligence: A Modern Approach. 4th ed. Pearson; 2020.

[7] Medexter Healthcare GmbH. Pipelines [Computer Software] 2026. Available from: https://pipelines.ardensyntax.com [cited 2026 March 20]

[8] Shortliffe EH. Computer-Based Medical Consultations: MYCIN. Elsevier; 1976.

[9] Allis LV, van der Meulen M, van den Herik HJ. Proof-number search. Artificial Intelligence. 1994;66(1):91-124. doi: 10.1016/0004-3702(94)90004-3

[10] Hopgood A. The State of Artificial Intelligence. Advances in Computers. 2005 12;65:1-75. doi: 10.1016/S0065-2458(05)65001-2