Synthesising Bayesian Network Models for Clinical Decision Support from Rule-Based Logic

Georgia Sowerby[1], Ol'Tunde Ashaolu[2], Radu Calinescu[1] and Stephen Connor[3]

[1]Department of Computer Science, University of York, York, United Kingdom
[2]York and Scarborough Teaching Hospitals NHS Foundation Trust, Scarborough, United Kingdom
[3]Department of Mathematics, University of York, York, United Kingdom

Abstract

We introduce a method for translating DAISY, a rule-based clinical decision support system, into per-malady Bayesian network models that incorporate uncertainty while preserving the underlying rule-based logic. DAISY gathers four categories of medically relevant patient data (demography, anatomical, subjective, and objective) and uses a decision-support algorithm to generate a set of potential assessments. We map each of DAISY's clinical variables to corresponding nodes in a Bayesian network, using directed edges to reflect the original "if-then" rules. Next, we apply parameter learning to estimate the conditional probability tables. To evaluate our method, we generated synthetic data that adheres to the DAISY rules and reflects realistic distributions based on clinical and population-level data. When applied to DAISY's knowledge base of over 200 maladies, this approach could result in substantially more informative triage reports, for example by ordering the list of potential assessments according to their calculated posterior probabilities.

Keywords

Clinical Decision Support System (CDSS), Bayesian network, Rule Based Expert System (RBES), Synthetic Data

1. Introduction

Emergency departments (EDs) worldwide face growing challenges due to increasing demand and limited resources. In the United Kingdom, the National Health Service Constitution handbook pledges a maximum four-hour ED waiting time, with an operational standard that 95% of patients are admitted, transferred, or discharged within this timeframe [1]. Yet in 2022-23, NHS Digital reported that over 25 million people attended EDs, with approximately 30% waiting longer than four hours to receive care [2]. This highlights strain on the system and raises concerns about timely and effective care. Furthermore, British Medical Association data reveal a significant shortage of doctors in England and many vacant positions [3]. These challenges are not unique to the UK: the World Health Organisation has reported that many European countries are facing "substantial shortages and gaps" [4]. The culmination of these pressures results in a challenging and high-stress working environment for ED staff, many of whom work long hours and report increasing levels of job dissatisfaction [5].

A key component in managing ED patient flow is *triage*, the *"clinical process to prioritise patients, completed before a full assessment to support effective management of demand and flow, identifying time critical requirements for patients"* [6]. This process typically comprises five stages [7]:

1. **Reception:** Administrative staff gather preliminary information by observing the patient and listening to their concerns. Patients who need immediate care can be escalated at this stage; otherwise, they progress to the next.
2. **History and Symptoms:** A triage clinician collects detailed patient information, asking questions about their medical history and current symptoms.

RuleML+RR'25: Companion Proceedings of the 9th International Joint Conference on Rules and Reasoning, September 22–24, 2025, Istanbul, Turkiye

georgia.sowerby@york.ac.uk (G. Sowerby); tunde.ashaolu@nhs.net (O. Ashaolu); radu.calinescu@york.ac.uk (R. Calinescu); stephen.connor@york.ac.uk (S. Connor)
0000-0002-6145-8336 (G. Sowerby); 0000-0002-7829-4884 (O. Ashaolu); 0000-0002-2678-9260 (R. Calinescu); 0000-0002-9785-2159 (S. Connor)

3. **Vital Sign Measurement:** Often occurring in tandem with the previous stage, the clinician measures and records the patient's vital parameters (e.g., temperature, heart rate, respiratory rate, blood pressure, oxygen saturation).
4. **Initial Assessment:** The clinician analyses the collected data to determine the patient's triage score and suggest potential assessments. This analysis can lead to various actions, such as escalating the case, returning the patient to the waiting room, or transferring them to a different department.
5. **Senior Clinician Review:** A senior clinician reviews the collected patient information and any potential assessments suggested by the clinician from Stage 4, then conducts a physical examination. They then decide on a treatment plan, which might include treating and discharging, referring for further investigations, admitting, or transferring to another facility.

This process is vulnerable to inconsistency, especially during busy periods. Clinical decision support systems (CDSSs) offer one possible solution to improve triage consistency and support overloaded clinicians by helping doctors make fair, evidence-based decisions regarding patient care. CDSSs are classified into knowledge-based or non-knowledge-based [8]. Knowledge-based systems use explicit rules, usually formulated as *if-then* statements, to represent medical knowledge. In contrast, non-knowledge based systems use machine learning (ML) techniques to identify patterns from large datasets. Although increasingly common in research, the real-world use of non-knowledge based CDSS's is limited due to concerns including explainability and limited access to high-quality data [8]. Knowledge-based systems are often preferred in clinical environments because their recommendations can be traced to clear, interpretable rules. However, these systems are deterministic and do not model uncertainty, which can limit their flexibility when information is incomplete or ambiguous.

One prototypical example of such systems is DAISY [9, 10, 7, 11], a knowledge-based CDSS developed to support ED triage by automating stages 2 through 4 of the process. It uses a rule-based architecture that prioritises transparency and explainability, but it does not model uncertainty or provide probability-ranked outputs.

In this paper, we present a method for converting DAISY into per-malady Bayesian network models, retaining its rule-based logic while enabling probabilistic reasoning. This enhancement allows the system to acknowledge uncertainty and generate probability distributions for possible assessments, rather than outputting simple binary decisions, potentially improving the informativeness of triage reports.

2. Related Work

Many clinical decision support systems have been proposed over the years. One of the earliest and most influential was INTERNIST-I, developed in the 1970s to address the growing complexity of internal medicine [12]. At its peak, the INTERNIST-I knowledge base included 572 diagnoses and over 4000 patient findings, with more than 4000 rules [13]. The rules were written by medical experts using textbooks and their own experience. However, the system had several limitations: it treated users as passive, required highly specific terminology for input, and generated consultations that lasted up to 75 minutes, which made it impractical for fast-paced clinical settings such as the ED.

To address these issues, Quick Medical Reference (QMR) was developed in the early 1980s. Like INTERNIST-I, QMR used the same knowledge base but acted more as an interactive information tool. It allowed users to input findings and receive feedback, supported by a completer feature to improve usability. Around the same time, ILIAD was introduced. While it began as a deterministic system, it later adopted a Bayesian network (BN) formalism. The resulting model included 11,406 nodes, with some structures extending to 36 levels and common findings shared by up to 62 parent nodes [14].

The transition to BNs in ILIAD reflected their advantages in handling uncertainty, a key challenge in medical reasoning [15, 16]. BNs also provide a clear graphical structure and can model causal relationships, making them well-suited to safety-critical applications like ED triage. In a comparative

study based in the ED, ILIAD outperformed QMR, providing correct diagnoses in 72% of cases versus QMR's 52%, although both systems generated long differential lists [17].

Machine learning has been used in emergency medicine applications such as predicting hospital admission [18, 19], workflow optimisation [20, 21], critical care [22, 23], and specific conditions such as sepsis or stroke [24, 25]. However, the lack of explainability in many of these models raises concerns [26, 8]. To address these concerns, recent research has focused on translating rule-based expert systems into Bayesian networks (e.g., [27, 28]). Our work builds on these developments by converting the DAISY triage system into a Bayesian network. DAISY was selected as the foundation for this research because it is actively maintained and specifically tailored to emergency triage. As members of the DAISY team, we have direct access to the codebase and to the clinical collaborators involved in its development. This makes it a more suitable platform than larger but less accessible systems such as INTERNIST-I.

3. Preliminaries

DAISY [9, 10, 7] gathers four categories of medically relevant information for an ED patient being triaged:

- **Demography:** The patient's medical history.
- **Anatomical:** Specific parts of the body affected.
- **Subjective:** Symptoms reported directly by the patient.
- **Objective:** The patient's vital signs.

Each category includes a set of variables that collectively describe the patient's clinical presentation in a structured manner. For example, the Demography category includes variables such as age, sex, and recent trauma. The Anatomical category includes variables that identify affected body parts such as head and chest. The Subjective category includes variables representing self-reported symptoms including pain and nausea. The Objective category includes variables for vital signs such as pulse rate and temperature.

The values of the variables within the Demography, Anatomical, and Subjective categories are obtained by asking the patient questions, which they answer via a touch-screen interface. The Objective category is obtained by instructing the patient to use medical devices in the room to measure their vital signs.

This information is processed by dAvInci, the rule-based expert system at the core of the DAISY project. It uses over 200 doctor-specified rules to output the patient's triage score and a set of potential assessments, suggested investigations, treatments and referrals. A triage report, which contains the information gathered by DAISY alongside the output from dAvInci, is made available to the patient's doctor.

DAISY is not merely conceptual. It has been implemented in a clinical setting and is currently undergoing a feasibility study at Scarborough Hospital as part of a registered clinical trial [11]. The study will involve 100 patients and will evaluate DAISY's acceptability, consultation duration, patient engagement, and clinical concordance compared to standard triage. These measures will provide insight into how well DAISY fits into real-world practice and how its outputs align with clinician assessments. The study will also generate real-world triage data to support the evaluation of the Bayesian network approach developed in this research.

To illustrate how DAISY's rules are expressed and how they can be translated into a Bayesian network, we introduce a simple example: the rule for SIRS; Meningitis[1]. This shows how demographic information, reported symptoms, and objective vital signs are combined within DAISY to trigger a malady assessment, and it will be used as a running example throughout the paper to demonstrate and explain our methodology.

[1]We use the label 'SIRS; Meningitis' to refer to cases where systemic inflammatory response syndrome (SIRS) is present in conjunction with or as a result of meningitis.

Example: SIRS; Meningitis

As a concrete example, consider the rule for identifying possible SIRS; Meningitis. The rule is triggered if the following conditions are met:

- **Demography** The patient has no history of recent physical trauma
- **Anatomy and Subjective** The patient reports a problem with their head and is bothered by bright lights (photophobia)
- **Objective** At least two of the following are true: abnormal temperature (e.g., low or high), elevated respiratory rate, elevated pulse rate

This rule can be formalised as:

$$
\begin{aligned}
&\texttt{Demography\_RecentTrauma} = \texttt{no} \\
&\wedge\ \texttt{Head\_BotheredByBrightLights} = \texttt{yes} \\
&\wedge\ \Big((\texttt{Objective\_Temperature} \neq \texttt{normal} \wedge \texttt{Objective\_RespiratoryRate} = \texttt{high}) \\
&\quad \vee\ (\texttt{Objective\_Temperature} \neq \texttt{normal} \wedge \texttt{Objective\_PulseRate} = \texttt{high}) \\
&\quad \vee\ (\texttt{Objective\_RespiratoryRate} = \texttt{high} \wedge \texttt{Objective\_PulseRate} = \texttt{high}) \Big) \\
&\implies \texttt{Malady\_SIRSMeningitis} = \texttt{yes}
\end{aligned}
$$

Our goal is to represent this logic within a Bayesian network whose probabilities can be used to relax the strict binary logic of these rules. Rather than requiring at least two abnormal vital signs to consider SIRS; Meningitis as a potential assessment, its probability increases with each abnormal vital sign. When more evidence is present, the probability changes accordingly, providing a more flexible, realistic way to support clinical decision-making under uncertainty.

4. Methodology

To translate DAISY's rule-based logic into a probabilistic model, we constructed separate Bayesian network models for each malady. We then generated synthetic training data and applied parameter learning to estimate their conditional probabilities. Each step of this process is illustrated using the running example of SIRS; Meningitis, introduced in Section 3.

4.1. Bayesian Network Construction

A separate Bayesian network is constructed for each malady represented in the DAISY knowledge base. Within each network, all variables from the DAISY categories Demography, Objective, and Malady are mapped to corresponding nodes. In DAISY, symptoms are always recorded with an associated anatomical variable, and these are combined in the Bayesian network into a single node type, Anatomy_Subjective. DAISY also includes a special Anatomy variable, 'General', to indicate that a symptom is not localised to a specific anatomical location (e.g., nausea).

Directed edges are added based on predefined causal assumptions representing clinical knowledge:

- **Demography → Malady** A patient's demographic and medical history can influence the probability of developing certain conditions.
- **Malady → Anatomy_Subjective, Objective** Once present, a malady is expected to cause symptoms and physiological changes.

Continuous variables, such as temperature and respiratory rate, are discretised into clinically meaningful intervals based on thresholds defined in DAISY's existing rule base. Although not reported here, we are also experimenting with representing these variables as continuous nodes within the Bayesian network to allow for more flexible modelling.

The networks are implemented in Bayes Server [29], with their structure defined using the Java API within R. The structure of the DAISY rule base permitted the majority of the network topology to be generated programmatically in a systematic and efficient manner.

4.2. Synthetic Data Generation

Due to the absence of suitable datasets, we generated synthetic data for 1000 patients using an R script, splitting the dataset into 700 cases for training and 300 for testing. The distributions from which the variables in the network were sampled were designed to reflect realistic emergency department conditions, while adhering to DAISY's original rule constraints. For the running example of SIRS; Meningitis, the variables and their associated distributions are shown in Table 1. These distributions were specified to provide sufficient training data while preserving clinically meaningful conditional relationships, such as a higher probability of photophobia when SIRS; Meningitis is present. The distribution parameters were set in consultation with clinical collaborators to approximate typical values for both normal and abnormal presentations.

Table 1
Variables used in the SIRS; Meningitis example

Variable	Type	Description	Distribution
Malady_SIRSMeningitis	Binary	Indicates whether the individual has SIRS; meningitis	Probability of "Yes" = 5%
Demography_RecentTrauma	Binary	Indicates whether the individual experienced recent physical trauma	When Malady_SIRSMeningitis = "Yes", probability of "Yes" = 15% When Malady_SIRSMeningitis = "No", probability of "Yes" = 40%
Head_BotheredByBrightLights	Binary	Indicates whether the individual reported sensitivity to bright lights	When Malady_SIRSMeningitis = "Yes", probability of "Yes" = 88% When Malady_SIRSMeningitis = "No", probability of "Yes" = 15%
Objective_Temperature	Continuous	Body temperature (°C)	When Malady_SIRSMeningitis = "Yes", Objective_Temperature is drawn randomly from either $N(35.5, 0.25)$ or $N(38.5, 0.35)$ with equal probability, bounded to [34, 41]. When Malady_SIRSMeningitis = "No", Objective_Temperature is drawn from $N(36.8, 0.7)$, bounded to [34, 41].
Objective_RespiratoryRate	Continuous	Respiratory rate (breaths per minute)	When Malady_SIRSMeningitis = "Yes", Objective_RespiratoryRate is drawn from $N(26, 7)$, bounded to [6, 35]. When Malady_SIRSMeningitis = "No", Objective_RespiratoryRate is drawn from $N(16, 7)$, bounded to [6, 35].
Objective_PulseRate	Continuous	Pulse rate (beats per minute)	When Malady_SIRSMeningitis = "Yes", Objective_PulseRate is drawn from $N(120, 35)$, bounded to [30, 180]. When Malady_SIRSMeningitis = "No", Objective_PulseRate is drawn from $N(75, 25)$, bounded to [30, 180].

4.3. Parameter Learning

The continuous variables in the synthetic dataset were discretised prior to training. Discretisation was carried out using interval-based state definitions, with thresholds chosen to reflect clinically meaningful

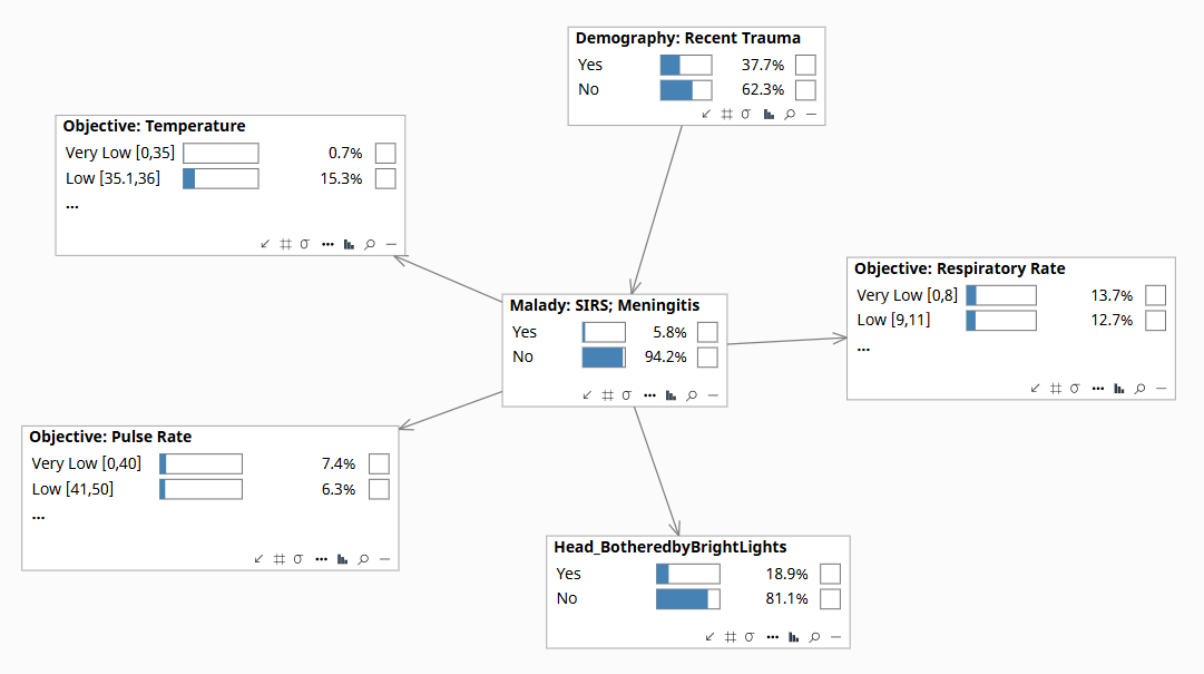

Figure 1: Trained Bayesian network for SIRS; Meningitis, with structure derived from DAISY's original rule logic.

ranges as defined in the DAISY rule base. For example, in the SIRS; Meningitis case, temperature was divided into five states (e.g., Very Low $[0, 35]$, Low $[35.1, 36]$, etc.), respiratory rate into five states, and pulse rate into six states.

With the Bayesian network structure defined according to DAISY's rule base, parameter learning was performed using the Relevance Tree inference algorithm provided by Bayes Server. This algorithm was chosen for its efficiency and compatibility with discrete and discretised variables.

5. Preliminary Results

We evaluated the Bayesian network for SIRS; Meningitis to assess whether the model produced clinically meaningful probability estimates and generalised well to unseen data.

5.1. Trained Bayesian network for SIRS; Meningitis

The structure of the trained network reflects DAISY's rule base, while the conditional probabilities were learned from the synthetic dataset described in Section 4.2. Figure 1 shows the resulting Bayesian network.

As expected, the model assigns higher probabilities of SIRS; Meningitis when more supporting evidence is present. For example, if a patient has no recent trauma, is bothered by bright lights, and has a high temperature and high pulse rate (but a normal respiratory rate), the model estimates a 76.9% probability of SIRS; Meningitis. If the respiratory rate is also high, this probability increases to 93.5%.

This example highlights the model's ability to capture variation in clinical presentation. Although certain conditions are associated with certain patterns, patients with the same diagnosis may differ in their medical history, affected anatomy, reported symptoms, and vital signs.

5.2. Discussion

Figure 2 presents box plots of the predicted probabilities for SIRS; Meningitis, grouped by the actual class label ("Yes" or "No"). The top plot shows results on the training set, and the bottom plot shows

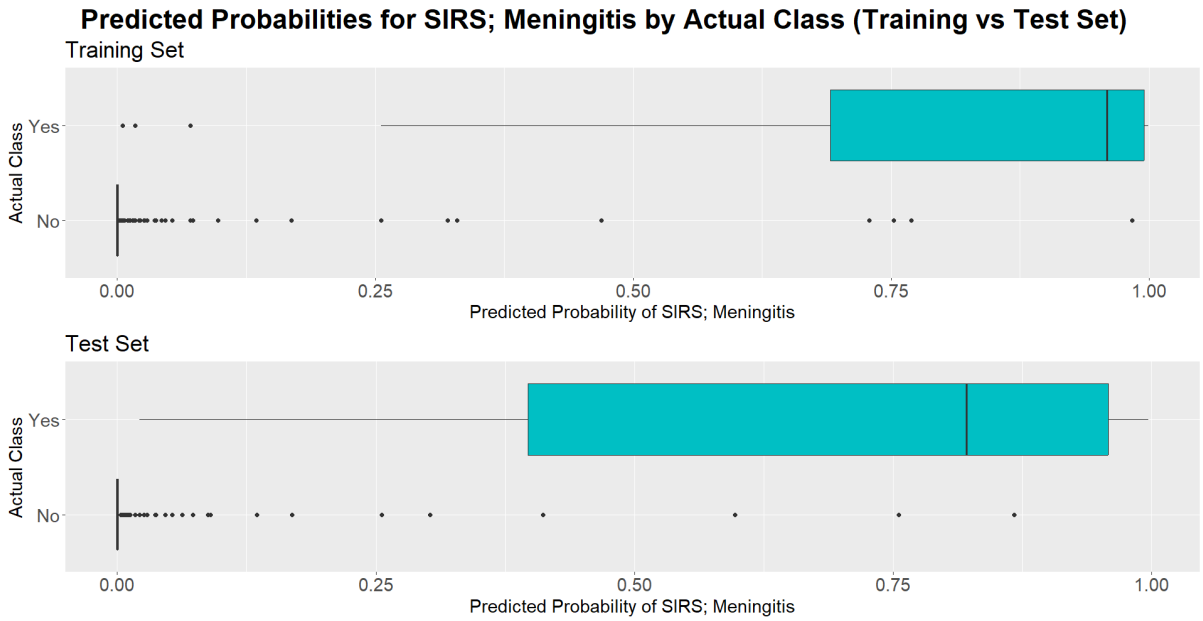

Figure 2: Predicted probabilities for SIRS; Meningitis by actual class. Top: Training Set; Bottom: Test Set

the test set.

In both cases, predicted probabilities are clearly separated between positive and negative classes. Patients in the "Yes" group consistently receive higher predicted probabilities than those in the "No" group, indicating that the model has learned to distinguish between the two classes effectively. The similarity between training and test results suggests that the model generalises well without signs of overfitting.

6. Conclusion

In this paper, we presented a method for translating the DAISY rule-based expert system into per-malady Bayesian networks, using the rule for SIRS; Meningitis as a running example. This approach preserves the explainability of rule-based systems while enabling the model to handle uncertainty and generate probability-ranked assessments rather than binary outputs.

Our preliminary results show that the trained Bayesian network produces clinically meaningful probability estimates, separates positive and negative cases effectively, and generalises well to unseen data.

Future work is in progress to extend this approach. We have constructed the structure of a Bayesian network for the entire DAISY knowledge base, covering over 200 maladies with approximately 300 nodes and more than 1500 links. This network has not yet been trained, but data from DAISY's ongoing feasibility study will provide real-world triage data for this purpose. Moving from per-malady models to a combined network will also allow us to investigate how co-morbidities can be represented, better reflecting the complexity of real-world clinical presentations.

Acknowledgments

This research is funded by a UKRI Trustworthy Autonomous Node in Resilience Doctoral Training Programme studentship.

Declaration on Generative AI

During the preparation of this work, the authors used GPT-4o in order to: Grammar and spelling check. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] Nuffield Trust, A&E waiting times, 2024. URL: https://www.nuffieldtrust.org.uk/resource/a-e-waiting-times.

[2] NHS Digital, Hospital Accident & Emergency Activity, 2022-23, 2023. URL: https://digital.nhs.uk/data-and-information/publications/statistical/hospital-accident--emergency-activity/2022-23.

[3] British Medical Association, NHS medical staffing data analysis, 2024. URL: https://www.bma.org.uk/advice-and-support/nhs-delivery-and-workforce/workforce/nhs-medical-staffing-data-analysis.

[4] World Health Organization, Ticking timebomb: Without immediate action, health and care workforce gaps in the European Region could spell disaster, 2022. URL: https://www.who.int/europe/news/item/14-09-2022-ticking-timebomb--without-immediate-action--health-and-care-workforce-gaps-in-the-european-region-could-spell-disaster.

[5] Sartini, Marina and Carbone, Alessio and others, Overcrowding in emergency department: causes, consequences, and solutions — a narrative review, Healthcare 10 (2022) 1625.

[6] NHS England, Guidance for emergency departments: initial assessment, n.d. URL: https://www.england.nhs.uk/guidance-for-emergency-departments-initial-assessment/.

[7] O. Ashaolu, W. Lyons, et al., Autonomous Emergency Triage Support System, in: CSCI'23, 2023, pp. 1332–1337. doi:10.1109/CSCI62032.2023.00220.

[8] R. T. Sutton, D. Pincock, et al., An overview of clinical decision support systems: benefits, risks, and strategies for success, NPJ Digital Medicine 3 (2020) 17.

[9] University of York, The DAISY project: Diagnostic AI System for Robot-Assisted A&E Triage, 2023. URL: https://www.cs.york.ac.uk/research/projects/daisy-project/.

[10] B. A. Townsend, K. L. Plant, et al., Medical practitioner perspectives on AI in emergency triage, Frontiers in Digital Health 5 (2023) 1297073.

[11] York Teaching Hospitals NHS Foundation Trust, Daisy - diagnostic ai system for robotic and automated triage and assessment (daisy), ClinicalTrials.gov ID NCT06571838, 2025. URL: https://clinicaltrials.gov/study/NCT06571838.

[12] J. D. Myers, The Background of INTERNIST I and QMR, in: ACM Conference on History of Medical Informatics, 1987, pp. 195–197. doi:10.1145/41526.41543.

[13] R. A. Miller, M. A. McNeil, et al., The INTERNIST-1/QUICK MEDICAL REFERENCE Project—Status Report, Western Journal of Medicine 145 (1986) 816.

[14] Y.-C. Li, et al., Automated transformation of probabilistic knowledge for a medical diagnostic system, Proc Annu Symp Comput Appl Med Care (1994) 765–769.

[15] E. Kyrimi, K. Dube, et al., Bayesian networks in healthcare: What is preventing their adoption?, Artificial Intelligence in Medicine 116 (2021) 102079.

[16] L. Uusitalo, Advantages and challenges of Bayesian networks in environmental modelling, Ecological Modelling 203 (2007) 312–318.

[17] M. A. Graber, D. VanScoy, How well does decision support software perform in the emergency department?, Emergency Medical Journal (2003).

[18] G. Feretzakis, G. Karlis, et al., Using machine learning techniques to predict hospital admission at the emergency department, The Journal of Critical Care Medicine 8 (2022) 107–116.

[19] W. S. Hong, A. D. Haimovich, et al., Predicting hospital admission at emergency department triage using machine learning, PLOS ONE 13 (2018) e0201016.

[20] M. Moustapha Mbaye, F. MH Abu Salem, et al., A new machine learning workflow to create an optimal waiting list in hospitals, in: ICMHI 2023, 2023, pp. 159–163.
[21] J. W. Joseph, E. L. Leventhal, et al., Machine learning methods for predicting patient-level emergency department workload, The Journal of Emergency Medicine 64 (2023) 83–92.
[22] Y. Liu, J. Gao, et al., Development and validation of a practical machine-learning triage algorithm for the detection of patients in need of critical care in the emergency department, Scientific Reports 11 (2021) 24044.
[23] Y. Raita, T. Goto, et al., Emergency department triage prediction of clinical outcomes using machine learning models, Critical Care 23 (2019) 1–13.
[24] T. M. Sullivan, Z. P. Milestone, et al., Development and validation of a Bayesian belief network predicting the probability of blood transfusion after pediatric injury, Journal of Trauma and Acute Care Surgery 94 (2023) 304–311.
[25] S.-F. Sung, L.-C. Hung, et al., Developing a stroke alert trigger for clinical decision support at emergency triage using machine learning, International Journal of Medical Informatics 152 (2021) 104505.
[26] H. Chang, W. C. Cha, Artificial intelligence decision points in an emergency department, Clinical and Experimental Emergency Medicine 9 (2022) 165.
[27] M. Korver, P. J. Lucas, et al., Converting a rule-based expert system into a belief network, Medical Informatics 18 (1993) 219–241.
[28] S. Thirumuruganathan, M. Huber, et al., Building Bayesian Network based expert systems from rules, in: IEEE SMC 2011, 2011, pp. 3002–3008.
[29] Bayes Server Ltd., Bayes Server, 2025. URL: https://www.bayesserver.com/.