Predicting the causative pathogen among children with pneumonia using a causal Bayesian network

Yue Wu[1*,2], Steven Mascaro[3,4], Mejbah Bhuiyan[2], Parveen Fathima[1], Ariel O. Mace[2,5,6], Mark P. Nicol[7], Peter Richmond[2,5,8], Lea-Ann Kirkham[2], Michael Dymock[2], David A. Foley[2,9], Charlie McLeod[2,10], Meredith L. Borland[8,11], Andrew Martin[5], Phoebe C.M. Williams[1,12,13], Julie A. Marsh[2], Tom Snelling[1,2,12,14,15], Christopher C. Blyth[2,8,9,10]

1 Sydney School of Public Health, University of Sydney, Camperdown, NSW 2006, Australia
2 Wesfarmers Centre of Vaccines and Infectious Diseases, Telethon Kids Institute, University of Western Australia, Nedlands WA 6009, Australia
3 Bayesian Intelligence Pty Ltd, Upwey VIC 3158, Australia
4 Faculty of Information Technology, Monash University, Clayton VIC 3800, Australia
5 Department of General Paediaitrics, Perth Children's Hospital, Nedlands WA 6009, Australia
6 Department of Paediatrics, Fiona Stanley Hospital, Western Australia, Australia
7 School of Biomedical Sciences, University of Western Australia
8 School of Medicine, University of Western Australia
9 Microbiology, PathWest Laboratory Medicine QEII Medical Centre, Perth, WA 6009, Western Australia
10 Infectious Diseases Department, Perth Children's Hospital, Nedlands, WA 6009, Australia
11 Emergency Department, Perth Children's Hospital, Nedlands WA 6009, Australia
12 Sydney Children's Hospital Network, NSW 2031, Australia
13 School of Women's and Children's Health, The University of New South Wales, NSW 2052, Australia
14 School of Public Health, Curtin University, Bentley WA 6102, Australia
15 Menzies School of Health Research, Charles Darwin University, Darwin NT 0815, Australia

* yue.wu1@sydney.edu.au

Abstract

Background

Pneumonia remains a leading cause of hospitalization and death among young children worldwide, and the diagnostic challenge of differentiating bacterial from non-bacterial pneumonia is the main driver of antibiotic use for treating pneumonia in children. Causal Bayesian networks (BNs) serve as powerful tools for this problem as they provide clear maps of probabilistic relationships between variables and produce results in an explainable way by incoporating both domain expert knowledge and numerical data.

Methods

We used domain expert knowledge and data in combination and iteratively, to construct, parameterise and validate a causal BN to predict causative pathogens for childhood pneumonia. Expert knowledge elicitation occurred through a series of group workshops, surveys and one-on-one meetings involving 6-8 experts from diverse domain areas. The model performance was evaluated based on both quantitative metrics (area under the receiver-operator curve (AUROC) and log loss) and qualitative expert validation. Sensitivity analyses were conducted to investigate how the target output is influenced by varying key assumptions of particular high degree of uncertainty around data or domain expert knowledge.

Results

Designed to apply to a cohort of children with X-ray confirmed pneumonia who presented to a tertiary paediatric hospital in Australia, the resulting BN offers explainable and quantitative predictions on a range of variables of interest, including the diagnosis of bacterial pneumonia, detection of respiratory pathogens in the nasopharynx, and the clinical phenotype of a pneumonia episode. Satisfactory numeric performance has been achieved including an AUROC of 0.8 in predicting the clinical diagnosis of bacterial pneumonia. Three commonly encountered scenarios were presented to demonstrate the potential usefulness of the BN outputs in various clinical pictures.

Conclusions

To our knowledge, this is the first causal model developed to help determine the causative pathogen for paediatric pneumonia. It can be utilized to derive recommendations to support more directed and judicious use of antimicrobials for relevant cohorts. The BN needs further validation before it can be clinically implemented. Our model framework and the methodological approach can be adapted beyond our context to broad respiratory infections and geographical and healthcare settings.

Keywords: paediatric pneumonia, Bayesian network, causal model, expert elicitation, clinical decision support

1 Introduction

Pneumonia, infection of the lower airways, remains a leading cause of hospitalisation and death among young children worldwide [1, 2]. In recent decades, marked reductions in the burden of disease have occurred thanks to public health strategies including immunisation targeting *Haemophilus influenzae* type b, *Streptococcus pneumoniae*, and influenza virus [3, 4], however substantial morbidity and mortality remains. These interventions has led to a decline in the contribution of bacterial pathogens in particular, challenging the appropriateness of routine empiric antibiotic use for treating community-acquired pneumonia (CAP) [5]. Viruses are increasingly considered a major cause of paediatric pneumonia especially in high income settings, and children with viral infection are unlikely to benefit from antibiotic therapy [6]; rather, unnecessary antibiotic use may lead to otherwise avoidable side effects, and be a driver of antimicrobial resistance [7, 8]. Despite this, antibiotics are still widely prescribed for children with pneumonia [5] possibly due to the difficulty in excluding bacterial infection as well as clinician concern regarding the potential consequences of

under-treatment of bacterial pneumonia whether caused by 'typical' bacterial pathogens such as *Streptococcus pneumoniae* (pneumococcus), atypical bacteria such as *Mycoplasma pneumoniae* (mycoplasma), or bacterial super-infection following viral infection [9].

The aetiology of CAP in children has been studied extensively [10]. A number of clinical and laboratory factors have been shown to be associated with pneumonia in sick children [11–13], and many of these features (including cough, increased respiratory effort and raised inflammatory markers) are common in both bacterial and non-bacterial pneumonia. However, no single perfect predictor exists that is sufficiently sensitive or specific to reliably differentiate bacterial (including bacterial-viral co-infection) from non-bacterial pneumonia [8]. Also, the performance of possible predictors may vary by context, e.g., in early versus late stage illness, or where the underlying prevalence of a pathogen varies [14]. The contribution of bacteria and viruses to pneumonia has been reported to vary by age, season, vaccine coverage, socioeconomic status, and across countries [6, 15, 16], although a recent review found evidence that viruses associated for a similar proportion of pneumonias across settings [17]. Isolation of pathogens directly from the lower respiratory tract is highly specific and may be sensitive, but sampling is typically invasive, challenging in children, and rarely indicated clinically and consequently rarely done [18]. Although less specific, nasopharyngeal samples are frequently used for extrapolating pneumonia aetiology because of ease of collection, however any pathogens detected in the nasopharynx might not be the actual cause of the pneumonia [19]. Due to the dynamic nature of disease epidemiology, non-specific clinical presentation and diagnostic limitations of laboratory tests, the timely and accurate identification of pathogens causing pneumonia remains an ongoing challenge.

Mathematical prediction models have been developed to aid clinical diagnosis, and existing approaches are usually based on quantifying associations between the target (e.g. bacterial pneumonia) and certain input variables in specified patient groups. Typical statistical regression methods are agnostic to underlying causal mechanisms and ignore complex interactions, such as those that occur between epidemiological, clinical, microbiological and immunological factors in pneumonia. Nevertheless, reasonable diagnostic performance has been achieved [20, 21]. However, the implementation and uptake of these models has been limited for several reasons. First, validation and clinical implementation of models are often challenged by missing data [22], such as biomarkers which might be useful predictors but may not be available in all cases to aid timely clinical decision-making [21]. Second, the models may not be transportable from the training context to other clinical contexts, because the predictive values of input variables are often driven by baseline prevalence [14], infection severity and testing techniques employed in the validated study population. Third, association-driven predictions that lack causal explanation may not be accepted by their intended end-users (i.e., clinicians) [23]. The clinical implementation of any model for decision support is largely dependent on trust which, in turn, depends on end-users understanding how the model works.

Bayesian networks (BNs) may offer a solution to the challenges of predictive model design and implementation into clinical practive [24]. BNs have been used to clarify complex medical problem domains [25], by facilitating probabilistic and causal reasoning using directed acyclic graphs (DAGs) [26, 27]. The graphical representation of BNs facilitates the translation of knowledge from domain experts and that acquired from clinical datasets into a causal inference framework, with quantitative relationships between variables. BNs provide probabilistic predictions that are transparent and explainable. The development of BNs relies on collaboration between domain experts and modellers to synthesize an understanding of the problem domain with the modelling technique, creating model outputs that are meaningful for clinical practice. BN model outputs can be applied to a variety of patient subgroups (such as specific age groups or geographic locations), and can robustly account for missing input variables like biomarker measurements from laboratory data. Our previous work has demonstrated how epidemiological, procedural and laboratory observations can be organised under a BN designed for predicting the causative pathogen in other childhood infections, and how such a model may be presented to aid clinical decision making [28, 29].

Here we present a causal BN model that depicts the pathophysiology of pneumonia in children, which could be used to help distinguish bacterial (including typical and atypical bacteria, and mixed viral-bacterial infections) from non-bacterial infections. The model was constructed, parameterised, and validated using both domain expert knowledge and data. We used data obtained from 230 children admitted to a tertiary hospital in Western Australia from 2015-2018 with pneumonia who underwent clinical, microbiological, immunological and radiological assessment [5, 30]. A diverse group of domain experts participated in the knowledge elicitation process. Based on numeric evaluation metrics, the outcome model showed stable and good performance in predicting clinical diagnosis of bacterial pneumonia as well as the detection of pathogens

in the nasopharynx. The model also produced clinically meaningful outputs as assessed by domain experts via evaluation workshops and surveys. To demonstrate the potential utility of the model, we apply it to three representative clinical scenarios and discuss how our model may be implemented as a tool for decision support to improve patient management and encourage judicious use of antibiotics for paediatric pneumonia.

2 Methods

A BN model consists of two components, 1) a DAG that qualitatively describes how variables (nodes) interact with each other, and 2) a joint conditional probability distribution that quantitatively specifies how changes in each parent variable probabilistically drive changes in their child variables. The arrows (arcs) of the DAG indicate the presence of a direct influence of predecessor (or parent) variables on their child nodes (nodes extending from other nodes). A causal BN is one in which the arcs represent influences that are causal; sometimes with minor exceptions for associations that are required, but are not a core part of the causal process being modelled. Developing a causal BN involves constructing the model structure (the DAG) that describes the problem domain of interest, and quantifying the probabilistic effect that parents have on their children in the model as a series of conditional probability tables (CPTs), i.e., parameterisation. In this project we used domain expert knowledge and data in combination and iteratively, to build a causal BN to predict causative pathogens for childhood pneumonia. Fig. 1 shows a schematic of the BN model development process.

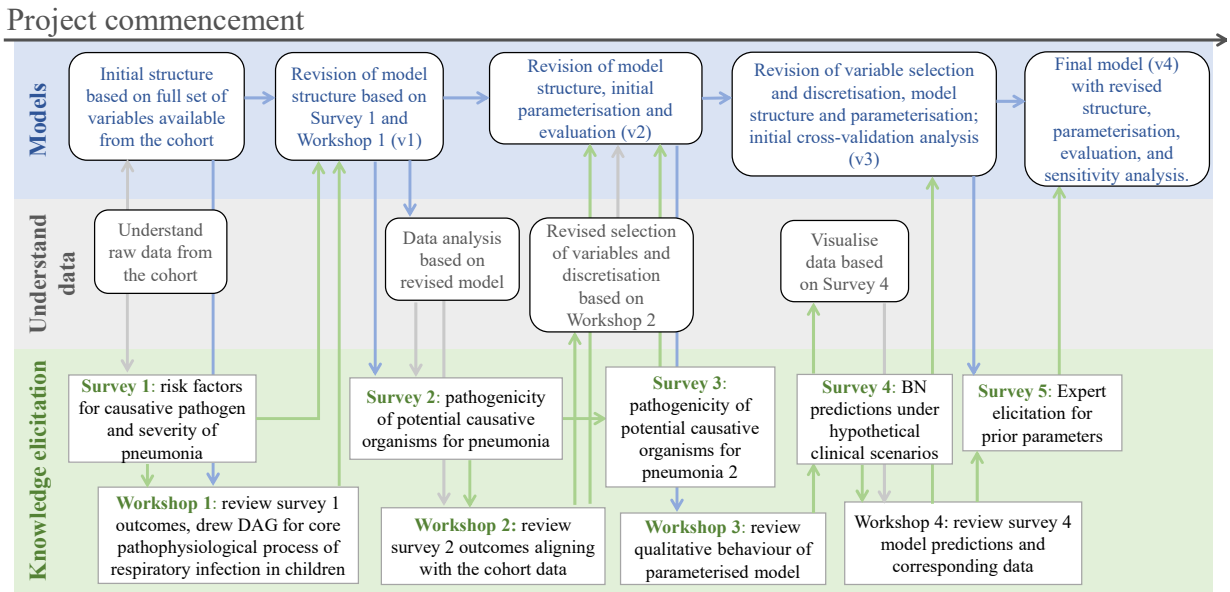

Fig 1. *Schematic of model development.* The top blue row presents the development of the BN from its initial form (left) to the fourth iteration (right), and how this iteratively interacts with the data analysis (middle grey row) and the knowledge elicitation (bottom green row). The four workshops were organised between December 2019 and August 2021.

2.1 Knowledge elicitation

Expert knowledge elicitation occurred through a series of group workshops, surveys and one-on-one meetings (Fig. 1) involving 6-8 domain experts (with backgrounds in infectious disease, clinical microbiology, emergency medicine, general paediatrics, immunology and epidemiology). Workshops and surveys covered the development of the model structure, elicitation of key model parameters, data interpretation, model

structure validation, and validation of the application use cases. One-on-one meetings were organised when more detailed discussion was considered necessary by the modellers, such as to consolidate understanding of a particular topic discussed during a workshop. As shown in Fig. 1, the theme of the workshop and survey was often driven by the evolving modelling progress—questions arising from data analyses, development of model structure and parameterisation. In Additional File 1, we provide a summary of survey questions.

2.2 BN structure and parameterisation

The final model structure was achieved through an iterative process of synthesising expert knowledge, data and learnings from earlier BN revisions (Fig. 1). Parameterisaton of the model was conducted using the expectation maximisation (EM) algorithm [31], with one of the following three methods applied to each node, depending on the quality of the data available for that node:

A ***data-driven*** approach was used to derive CPTs directly from the observed distribution of relevant variables in the data, when data for those variables was sufficient.

An ***expert-driven*** approach was used for conceptual variables that were not or could not be recorded in the data. In this approach, the CPT was fully specified by domain expert knowledge alone and not further updated with data. For example, **susceptibility to progression** describes the extent to which an infected child is at risk of progressing from mild to severe manifestations of pneumonia. In the model, this susceptibility is higher in younger children, in those with chronic respiratory disease, and in those who are immunocompromised.

A ***hybrid expert-data*** approach was used in which expert-driven CPTs were used as a Bayesian prior (i.e., a starting point) and subsequently updated using available data. This method was preferred for variables that were strongly associated with other data, but could not be directly observed or mapped to any individual variable (i.e., latent), as well as for observable variables where data were sparse for certain patient subgroups.

2.3 BN performance and evaluation

We evaluated the performance of the BN in two ways. Quantitatively, numerical evaluation ensured the model predictions were consistent with the observed data. While metrics that predict accuracy such as area under the receiver-operator curve (AUROC) have been widely used to summarise the performance of classifiers regarding their true positive/negative and false positive/negative rates, they can be misleading. These measures are more suitable for definite classifications rather than the probabilistic predictions that some models are able to provide. For probabilistic models like BNs, we can create a binary classification by applying a threshold. For example, if we choose a cut-off threshold of 0.5, the model predicted probabilities of 0.51 and 0.99 would be collapsed together; if we instead shifted the threshold to 0.55, then 0.51 and 0.01 would be collapsed together. AUROC mitigates this problem by evaluating all possible thresholds, but still rests on the expectation that outcomes will be dichotomised one way or another — when in fact end-users of the model will see the probability, not a deterministic prediction. For this reason, log loss (and similar probability-sensitive metrics) is considered more informative for evaluating models that generate probabilistic predictions [32]. Log loss is calculated, for a given case, as the negative log of the model's estimated probability of what was actually observed (i.e. $-log\{P(observed\ outcome)\}$) with a lower log loss indicating a better fit of the model to the observed data. The log loss metric rewards accurate probability estimates most highly. For example, if a positive blood culture is truly observed, a model predicted probability for positive blood culture of 0.51 yields a log loss of –log (0.51) = 0.29, while a model predicted probability of 0.99 yields a much lower (and hence better) log loss of 0.004, reflecting a reward for the greater confidence in a positive culture if correct, but punishing this greater confidence if incorrect. Accumulating log loss over a number of cases then rewards the most accurate probabilistic predictions.

Sensitivity analyses were conducted using variance-based sensitivity analysis (VBSA) [33,34], to investigate how the target variable of the model is influenced by varying the assumptions for a range of key variables, especially those affected by a high degree of uncertainty around data or domain expert knowledge.

The approaches above are unfortunately of limited value in evaluating the causal features of a model. For example, its common for two BNs to differ in their causal features but be identical in their statistical

properties [27, 35]; such models would be assessed as equivalently good by any of the above metrics. More commonly, major differences in causal features may only lead to minor differences in predictive performance. Thus, while predictive performance can help identify significant issues (such as major missing pathways or dependencies), other approaches must be used to evaluate the individual causal features of a BN. This can be achieved in a number of ways (e.g., comparing to literature, conducting experiments, etc.); here we performed this validation with experts, by walking through the DAG structure afresh at each workshop, and by soliciting feedback on the DAG structure via followup emails.

In addition, we performed qualitative expert validation of the model outputs, allowing us to assess if the model behaves in a way that meets clinician expectations and offers clinically informative predictions. We simulated three typical scenarios encountered in clinical practice, and presented BN predictions in both survey format (Survey 4) and during workshops (Workshops 3 and 4) for expert review.

Ethics approval was granted by the Child and Adolescent Health Service Human Research Ethics Committee (RGS2477). We used the GeNIe BN (`https://www.bayesfusion.com/`) software to elicit the BNs presented here, and Netica (`https://www.norsys.com/`) to develop and parameterise subsequent models. Data, evaluation and sensitivity analyses were conducted using R [36] and Python [37]. Throughout the paper, we label the name of a BN variable in **bold**, and its state name in *italics* where referred for the first time.

3 Results

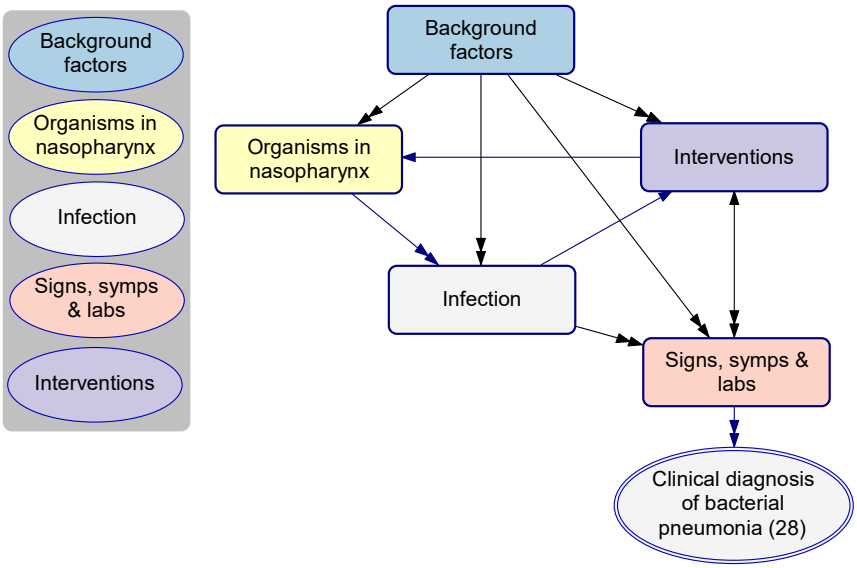

Fig 2. The high-level model structure of paediatric pneumonia (BN v4.5). Variables were organised into five highly interdependent groups (shown as rectangle boxes). In the centre of the diagram sits the core part of the model, infection of the respiratory tract (white box), which is predominantly influenced by the presence of a range of pneumonia-causing pathogens in a child's nasopharynx (yellow), as well as a number of background factors (blue). Infection manifests as clinical signs, symptoms and abnormal laboratory markers (salmon). Interventions (purple) refer to investigations and treatments which are relevant to the episode of pneumonia, and which may either occur before presentation (e.g., prior antibiotic use) potentially affecting variables like the culture result, or after presentation (e.g., supplemental oxygen) which may be affected by variables like the infection phenotype. Additional File 2 includes a BN dictionary which defines each variable.

Fig. 2 presents the high-level structure of the BN (v4.5), a model of paediatric pneumonia at the point of hospital presentation. We explicitly model the *existence* of infection (which is latent) separately from

its clinical diagnosis (white node) which may be operationally defined by a set of clinical or laboratory observations. A more detailed description of the model is provided in §3.2.

3.1 Summary of model variables

The model consists of 63 variables including 14 which are background factors, six pathogen-specific variables describing the presence of pathogens in the nasopharynx, nine variables representing infection or the diagnosis of infection, 24 variables representing specific signs and symptoms, six laboratory variables, and four intervention variables. Alongside the 53 observable variables, nine latent variables were created for modelling purposes; latent variables help to explicitly describe important underlying mechanisms that cannot be directly observed and thus captured by data. One variable (**current clinical phenotype (52)**) was introduced as a summary of the presenting clinical phenotypes of pneumonia. This variable was derived based on a separate clustering using the EM algorithm to determine a child's most probable pneumonia phenotype based on their clinical observations. The analysis resulted in two phenotypes, where *phenotype 1* represented a more severe type which increased the probability of all observable signs and symptoms compared to *phenotype 2*, and this information was added as an additional column in the cohort data and subsequently used as an observed variable. We discuss the role of this variable further below §3.2.3.

In the BN dictionary, we provide information on each model variable, including identifiers (ID and variable name), what it means and how it's mapped to data (description), how it's discretised in the BN (states), its parents in the BN, and how it's affected (mechanistically rather than statistically) by those parent nodes. The dictionary also notes the observational status (whether observable, latent or derived), and the parameterisation method used (data-driven, expert-driven or hybrid expert-data) for each variable. Fig. 3 shows a sample of the BN dictionary; the full dictionary is provided in Additional File 2.

ID	Variable name	Description	States	Parents in the BN	Relationship with parent nodes	Status	Parame-terisation	Category
1	**Age group**	Age group of study participant. In the model, we define each group as follow: Infant (<=2yo), …	*Infant, PreSchool, School*	None	Not applicable	Observable	Data-driven	Background factors
14	**Susceptibility to progression**	This describes the extent of the child to progress to more severe manifestation of pneumonia if infected.	*High, Low*	Age group, chronic respiratory disease, impaired immunity	The pneumonia is more likely to progress if the child's immune system is unable to clear the infection…	Latent	Expert-driven	Background factors
21	**Viral-like nasopharyngeal infection**	Replication of viral-like pathogens is occuring in the nasopharyngeal tissues.	*Present, Absent*	Organisms in nasopharyn (influenza, RSV, HMPV, parainfluenza, mycoplasma)	Presence of virus or mycoplasma in the nasopahrynx predisposes viral-like nasopharyngeal infection.	Latent	Expert-data	Infection
44	**Crackles**	Crackles/ crepitations as auscultatory finding recorded in the medical notes.	*Recorded, Unknown*	Current clinical phenotype, causative pathogen for pneumonia	Crackles refers abnormal lung sounds characterized by discontinuous clicking or …	Observable	Expert-data	Signs & symps
55	**C-reactive proteins**	The amount of C-reactive proteins (CRP) detected from blood.	*Above70, Btw30And70, Below30*	Current clinical phenotype, causative pathogen for pneumonia	Elevated level of CRP in blood can be driven by the systemic inflammatory response, which …	Observable	Expert-data	Labs

Fig 3. *Extract from the BN dictionary.* The full dictionary is provided in Additional File 2.

3.2 Model description

3.2.1 The epidemiological context

The model introduced three latent concepts (light blue) to simplify and describe the complex nature of pneumonia epidemiology in children; namely the **level of exposure (12)**, the **susceptibility to colonisation (13)**, and the **susceptibility to progression (14)** (Fig. 4). The susceptibility to colonisation summarises the level of a host's susceptibility to being colonised in the nasopharynx with 'typical' bacterial pathogens, which is higher in the younger **age group (1)**, in children with **impaired immunity (6)**, and when there is a **smoker in the household (3)**. The level of exposure refers to the host's exposure to transmissible pathogens. Greater exposure is associated with older age, more **childcare days per week (5)**, and *Indigenous* Australian **ethnicity (2)**, which may in turn be surrogates of frequent and/or close or prolonged social contact and therefore greater opportunity for transmission. The susceptibility to progression

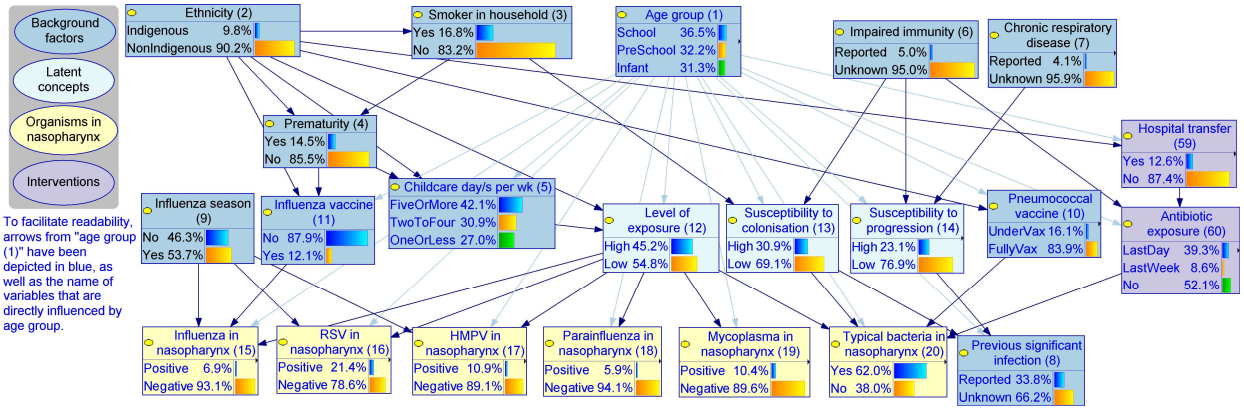

Fig 4. *The epidemiological context, BN v4.5.* Due to the number of arcs emanating from Age group, a lighter shade of blue has been used to make the graph easier to read. The detailed model structure and definition for each variable are provided in Additional File 2. Source model files can be accessed via Open Science Framework (`https://osf.io/m97vb/`).

describes the propensity of the host to pneumonia progression if infected, which is higher if the child is younger, is reported to have impaired immunity or has **chronic respiratory disease (7)**.

This BN explicitly describes the presence of six pathogens in the nasopharynx: influenza, respiratory syncytial virus (RSV), human metapneumovirus (HMPV), parainfluenza, mycoplasma and typical bacteria (e.g., pneumococcus) (Fig. 4, yellow nodes **(15) - (20)**). Various epidemiological risk factors interact resulting in varied prevalence of each pathogen in the nasopharynx, which predispose children to respiratory tract infection including pneumonia, leading to the core part of the model presented in Fig. 5. For example, while seasonality can influence the prevalence of many pathogens, in this model we use **influenza season (9)** (defined as June to September based on the expected pattern of influenza in the southern hemisphere in the pre-COVID-19 era) as a surrogate of the temporal increase in the prevalence of influenza, RSV and HMPV, with the model estimating increases of 220%, 173% and 177%, respectively. **Influenza vaccine (11)** reduces the probability that influenza is present in the nasopharynx by a modelled estimate of 67%. Recent **antibiotic exposure (60)** reduces the probability of typical bacteria being present in the nasopharynx from 75% to 49%.

3.2.2 The infection

Fig. 5 shows how the BN depicts infection, based on whether viruses or typical or atypical bacteria are involved in a particular pneumonia episode, and whether upper respiratory sites (nasopharynx and/or throat) are involved when pneumonia is present. Mycoplasma is an intracellular bacterial pathogen with similarities to viruses with respect to mechanism of acquisition and some clinical manifestations. For modelling purposes, we grouped it with the viruses (described as the *viral-like* group) while keeping all the other bacteria in the *typical bacterial* group. For viral-like pathogens, the BN assumes infection of the lung (**viral-like pneumonia (23)**) can either occur directly without observed upper respiratory tract infection (URTI), or it can occur following observed URTI including **viral-like nasopharyngeal infection (21)** and **viral-like throat infection (22)**. The exact probability of each virus causing a respiratory infection is influenced by its pathogenicity (parameters estimated via the hybrid expert-data approach). Viral infection of the nasopharynx or throat increases the probability of viral pneumonia. Unlike viral-like pathogens, typical bacteria frequently colonise in the nasopharynx (62% in the model), described by **Typical bacteria in nasopharynx (20)**, and spontaneously cause pneumonia (**typical bacterial pneumonia (24)**) at a much lower frequency than viruses. In addition, pneumococcal vaccination, the presence of viral-like infection, and the age of the child, can all directly influence the probability of typical bacterial pneumonia.

We introduced the **causative pathogen for pneumonia (25)** and **upper airway involvement (26)** variables to summarise the above. We defined the causative pathogen for pneumonia to be *typical bacterial* if **typical bacterial pneumonia (24)** is *present*, *viral-like* if **viral-like pneumonia (23)** is *present*

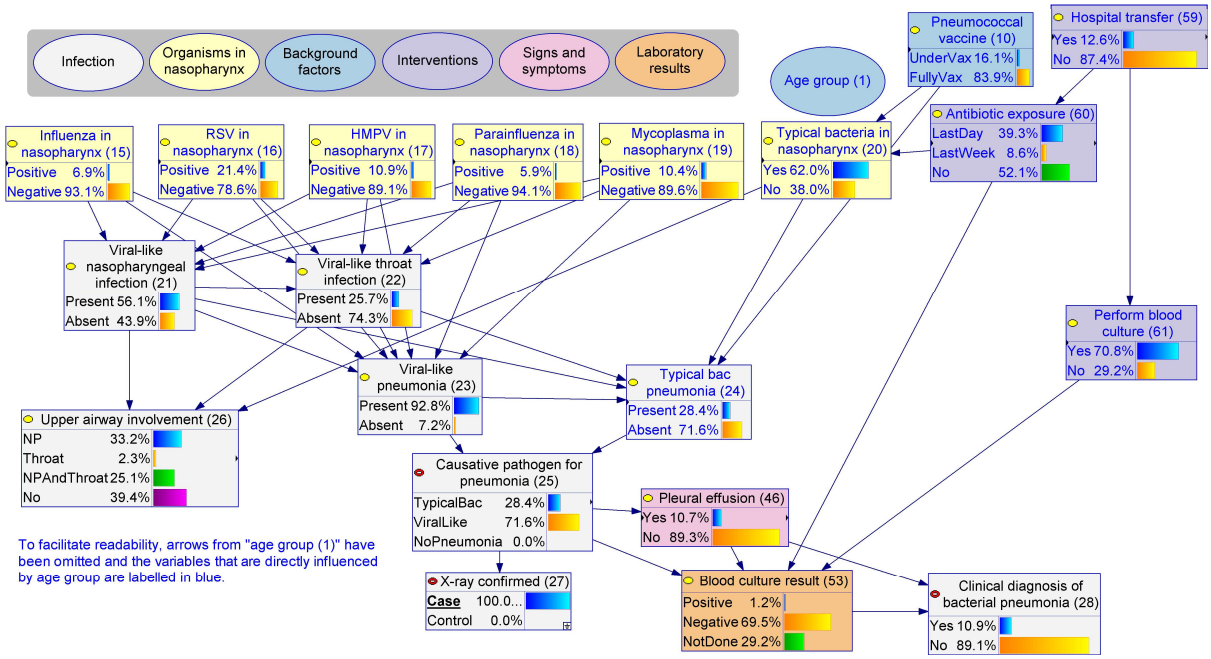

Fig 5. *The infection, BN v4.5.* The detailed model structure and definition for each variable are provided in Additional File 2. Source model files can be accessed via Open Science Framework (`https://osf.io/m97vb/`).

and typical bacterial pneumonia is *absent*, or *no pneumonia* if both the typical bacterial and viral-like pneumonia are *absent*. Note that the probability of *no pneumonia* should always be 0% in our model as it was parameterised using data from a cohort of children who all had **X-ray confirmed pneumonia (27)**. The operational definition of **clinical diagnosis of bacterial pneumonia (28)** in the cohort study was either having a **pleural effusion (46)** or a *positive* **blood culture result (53)**, comprising 11% of children in the cohort. As illustrated in Fig. 5, this definition is subject to potential measurement and selection biases introduced by prior antibiotic exposure, the decision to **perform blood culture (61)** and **hospital transfer (59)**.

The BN predicted that of all pneumonia presentations, 28.4% involved infection with typical bacterial pathogens with a mean of 25.2% based on the 10-fold cross validation (10-fold mean), while 92.8% (10-fold mean 94.6%) involve infection with viruses or viral-like pathogens. Among all pneumonia presentations involving infection with viruses or viral-like pathogens, 22.8% (10-fold mean 20.1%) were predicted to also involve infection with typical bacterial pathogens. From the decision-support perspective, it is important to differentiate infection with mycoplasma from viruses, and this can be addressed by combining information provided by two nodes, the **causative pathogen for pneumonia (25)** and **mycoplasma in nasopharynx (19)**. In other words, if the model predicts a high probability of infection with a *viral-like* pathogen and also a high probability of mycoplasma in nasopharynx, there is a high probability that mycoplasma is the causative pathogen (with possible implications for management).

3.2.3 The evidence

From a clinical perspective, a child's **causative pathogen for pneumonia (25)**, **current clinical phenotype (52)** and **upper airway involvement (26)** are not directly observed (i.e., they are latent), and can only be inferred from clinical observations and the results of laboratory investigations. From a causal perspective, these latent variables interact to give rise to the various types of evidence including clinical signs, symptoms, and laboratory results (Fig. 6). The potential causes of signs, symptoms, and laboratory results are rarely deterministic, singular or mutually exclusive. For example, **cough (29)** can be driven by upper airway irritation, inflammation, or mucus production (e.g. postnasal drip secondary to rhinorrhoea/rhinitis)

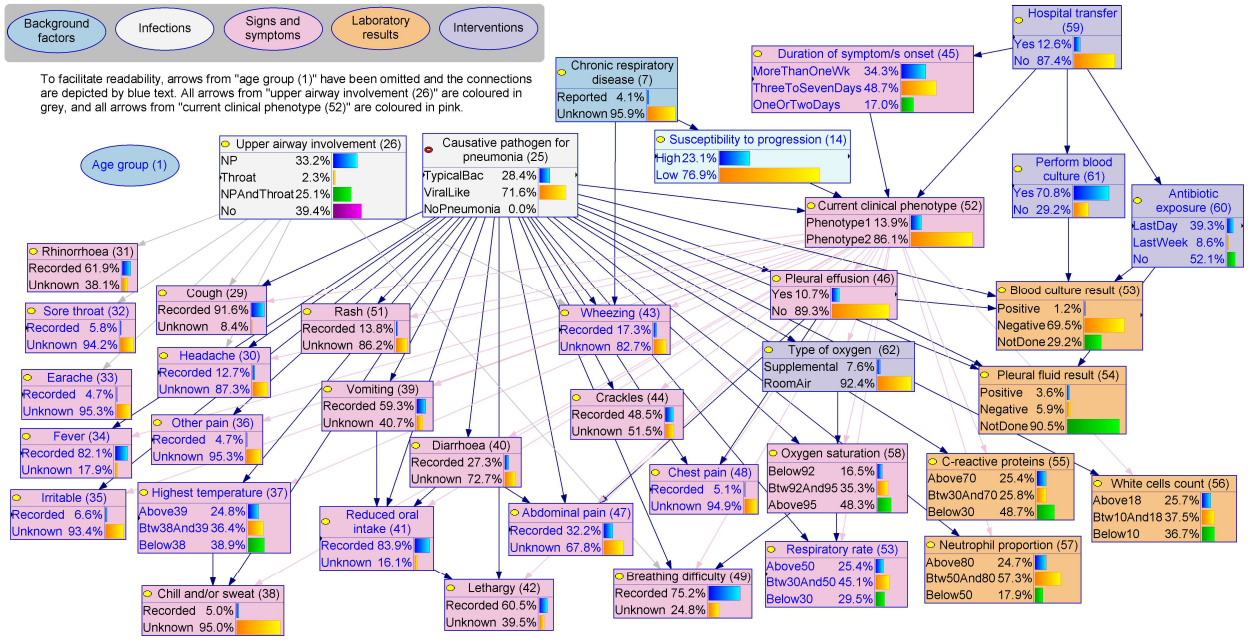

Fig 6. *The evidence, BN v4.5.* Causative pathogen for pneumonia, current clinical phenotype and upper airway involvement interact to give rise to the various types of evidence including clinical signs, symptoms, and laboratory results. The BN dictionary (Additional File 2) details how these latent variables may give rise to each type of evidence. Source model files can be accessed via Open Science Framework (`https://osf.io/m97vb/`).

or can be from pneumonia affecting the lower airways. Similarly, increased **respiratory rate (53)** can be a consequence of airway inflammation or low **oxygen saturation (58)** and can vary by the child's age.

Background (dark blue) and intervention (purple) factors may also influence certain types of evidence. For example, age is an important determinant of the report of pain (e.g., **headache (30)**, **abdominal pain (47)** and **sore throat (32)**) because children need to be old enough to communicate these symptoms. Recent **antibiotic exposure (60)** may reduce the probability of a successful blood culture (**blood culture result (53)**) even if a child is bacteremic. Knowing the **type of oxygen (62)** (*supplemental* or *room air*) is essential for interpreting oxygen saturation in relation to lung function. Symptoms and signs may also interact, e.g., **vomiting (39)** and **diarrhoea (40)** may lead to **reduced oral intake (41)** which could subsequently lead to **lethargy (42)**.

The classification of the current clinical phenotype of pneumonia is analogous to a count-based metric for disease severity, where those with *phenotype 1* had a median of 12 pneumonia-relevant signs and symptoms, compared to a median of 8 for *phenotype 2*. The inclusion of this variable allowed us to separate the differential influence of causative pathogen (typical bacterial or viral-like) from the severity of clinical manifestations conditional on the causative pathogen. For example, the overall probability of a child being *phenotype 1* (severe) was almost identical for pneumonia involving viruses and viral-like pathogens (14.0%) and those involving typical bacterial pathogens (13.6%). For a child who had symptom onset was within two days and been transferred from another hospital, the probability of *phenotype 1* was much higher for viral-like pneumonia (79.4%) than typical bacterial case (42.6%). In contrast, if the child presented to ED directly (not transferred), the overall probability of *phenotype 1* is lower (15.5%), however, typical bacterial pneumonia has a more severe phenotype (23.8%) than their viral-like counterparts (12.1%).

3.3 Model evaluation and sensitivity analyses

Fig. 7 presents the BN predicted probabilities of clinical diagnosis of bacterial pneumonia for every pneumonia episode for children in the cohort, and compares these against their final diagnosis (i.e., Bacterial Diagnosis:

Yes/No). Note that we use the recorded clinical diagnosis as the best estimate of the causative pathogen type, acknowledging that some of these diagnoses may be incorrect. So long as the number of incorrect diagnoses is not great, this will not affect our conclusions. The graphs represent four input scenarios ***(a-d)***, each providing more information to the model than its preceding scenario:

(a) provides the model with available basic background factors of the partipants; these can include age group, childcare days per week, chronic respiratory disease, ethnicity, influenza season, prematurity, smoker in the household, impaired immunity, influenza vaccine, pneumococcal vaccine, and previous significant infection.

(b) uses the information in ***(a)*** plus signs, symptoms and interventions if reported; these can include duration of symptom/s onset, rhinorrhoea, cough, respiratory rate, oxygen saturation, breathing difficulty, crackles, wheezing, (parent-reported) fever, highest temperature, chill, sweats, irritability, vomiting, diarrhoea, reduced oral intake, lethargy, chest pain, abdominal pain, other pain, earache, sore throat, headache, rash, (prior) antibiotic exposure, type of oxygen, hospital transfer, and (decision to) perform blood culture.

(c) uses information in ***(b)*** plus any detection (or no detection) of respiratory pathogens in the nasopharynx; these can include RSV, influenza, HMPV, parainfluenza, mycoplasma, and any typical bacteria.

(d) includes ***(c)*** plus all other available results; these can include C-reactive proteins (CRP), white cell count (WCC), and neutrophil proportion (of total WCC).

Graphs that show a strong overlap (such as in Fig. 7a) indicate the model is only able to achieve weak discrimination between bacterial diagnoses and others. However, as further information is provided (through scenarios (b), (c) and (d)), the discrimination improves, as evidenced by both the graphs and the improving AUROC scores. In contrast, the log score does not improve over the first 3 scenarios; coupled with the improving AUROC, this suggests the initial probabilistic prediction is already quite close to the true class. Log loss does improve in the final scenario, as does AUROC again and the greatest discrimination is visible in this graph. Table 1 additionally presents the evaluation metrics for other targets including pleural effusion, positive blood culture result, and the detection of viruses or mycoplasma in the nasopharynx.

Table 1. The performance of BN v4.5 in predicting a range of evaluation targets under different input scenarios, in the form of mean log loss and AUROC based on 10-fold cross-validation. Note for each target and scenario, the variables used to operationally define the evaluation target were removed from the input variables.

Evaluation target	**AUROC**				**Log loss**			
Scenario	***(a)***	***(b)***	***(c)***	***(d)***	***(a)***	***(b)***	***(c)***	***(d)***
Clinical diagnosis of bacterial pneumonia	0.58	0.70	0.71	0.80	0.36	0.36	0.35	0.31
Pleural effusion	0.58	0.73	0.73	0.81	0.35	0.33	0.33	0.30
Positive blood culture	0.73	0.74	0.80	0.91	0.67	0.08	0.08	0.06
RSV	0.74	0.75	NA	0.75	0.48	0.51	NA	0.52
Influenza	0.52	0.56	NA	0.56	0.29	0.28	NA	0.28
HMPV	0.53	0.54	NA	0.54	0.36	0.36	NA	0.36
Parainfluenza	0.45	0.49	NA	0.52	0.25	0.25	NA	0.24
Mycoplasma	0.67	0.73	NA	0.73	0.33	0.32	NA	0.32

We varied relevant parameters by 20% to investigate how various factors affect the predicted probability of typical bacterial pneumonia (Fig. 8). This target prediction is most sensitive to parameters describing the pathogenicity of the potential pathogens in causing pneumonia, ranging from 12% to 50%. Varying the prevalence of various pathogens in the nasopharynx caused the predicted probability of the involvement of

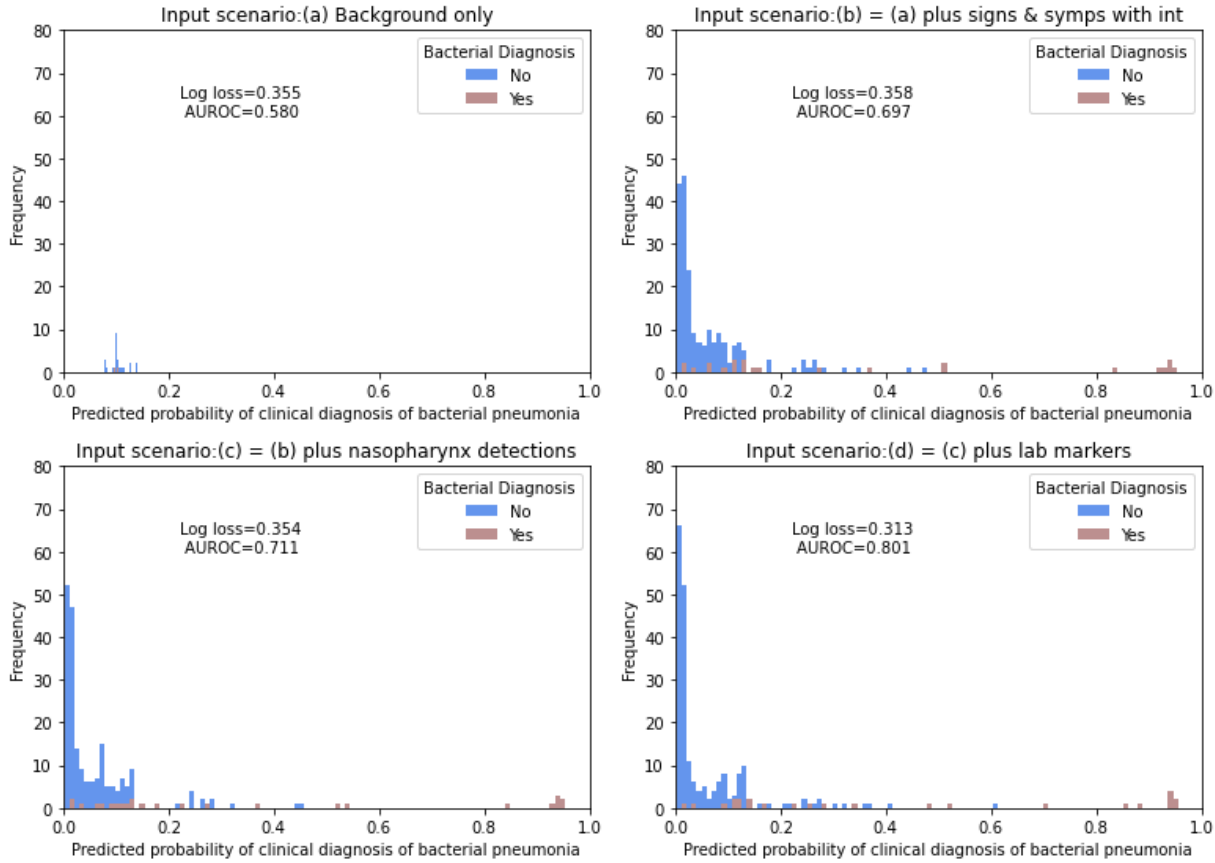

Fig 7. *The performance of BN v4.5.* Predicted probabilities of clinical diagnosis of bacterial pneumonia are compared with the final bacterial diagnosis for four input scenarios, with mean log loss and AUROC based on 10-fold cross-validation. For each scenario, available information from each participant was entered as input for the BN to predict the clinical diagnosis of bacterial pneumonia.

typical bacteria to vary from 22% to 30%. Varying the assumptions regarding the latent epidemiological concepts (e.g., level of exposure) had a small influence on the probability of involvement of typical bacteria (25% to 27%).

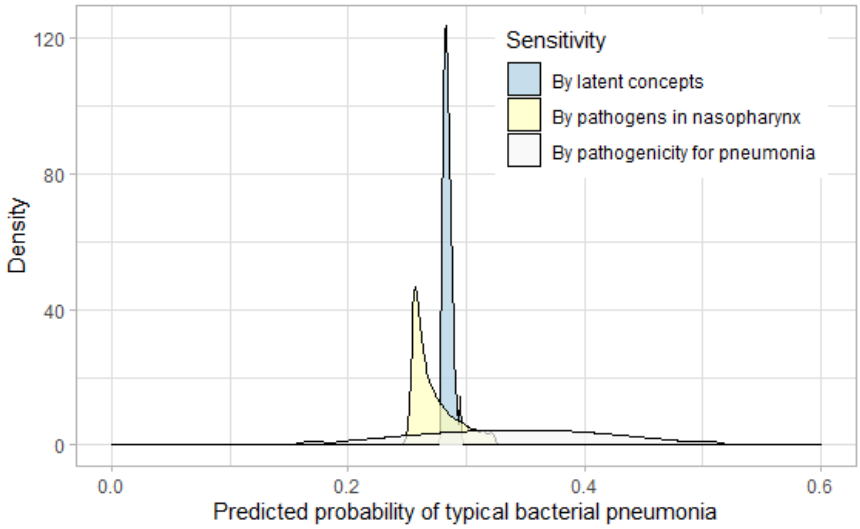

Fig 8. *Sensitivity analysis, BN v4.5.* Relevant parameters were varied by 20% to investigate how they affected the predicted probability of involvement of typical bacteria. This target prediction was most sensitive to the pathogenicity of potential pathogens in causing pneumonia (white), followed by the prevalence of various pathogens in nasopharynx (yellow), and was least sensitive to the latent epidemiological concepts (blue).

3.4 Demonstrative clinical scenarios

When serving as a clinical decision support tool, the available background factors, clinical observations and investigation results would be entered as model inputs for making inferences about the pathogen causing pneumonia in a particular episode. We presented three scenarios to demonstrate the potential use of this model in a clinical context (Fig. 9-11). The BN predicted probability of typical bacterial pneumonia is 28% for the cohort of children with X-ray confirmed pneumonia (labelled in the blue box at the top right corner). For each scenario, this prediction is updated, branching in multiple possible ways that are conditional on information that may become available over time along the child's clinical course. In parentheses we provide the change in the target prediction in relation to the cohort average; we use red text to indicate if the predicted probability of typical bacterial pneumonia is higher than the cohort average, and use green text to indicate predicted target probabilities which are more than 50% lower than the cohort average.

Fig. 9 presents a scenario in which the BN predicts a low probability (13%) of a typical bacterial cause of pneumonia for a 4 month old infant with X-ray confirmed pneumonia with moderate signs and symptoms (i.e., highest temperature of 38.5, runny nose and cough). The predicted probability remains low unless both the CRP and WCC are high (at least 100 and 12, respectively). In contrast, if instead a 9 year old (school-aged) child with X-ray confirmed pneumonia presents various signs and symptoms (Fig. 10), the probability of typical bacterial pneumonia is predicted to be higher and similar to the cohort average. The predicted probability of typical bacterial pneumonia is almost doubled if the child has no breathing difficulty, and a moderately high CRP (40) can be indicative for a bacterial cause. In other words, the significance of a raised CRP varies under various scenarios. The third scenario (Fig. 11) presents a 3.5 year old (pre-school aged) child with X-ray confirmed pneumonia and very high temperature during the influenza season. The probability of typical bacterial pneumonia increases if the child is under-vaccinated for pneumococcus, regardless of whether RSV is detected in their nasopharynx. If fully-vaccinated, the presence of wheezing will lead to a decrease in the predicted probability of typical bacterial pneumonia, and this prediction will drop to only 4% if RSV is detected in the child's nasopharynx.

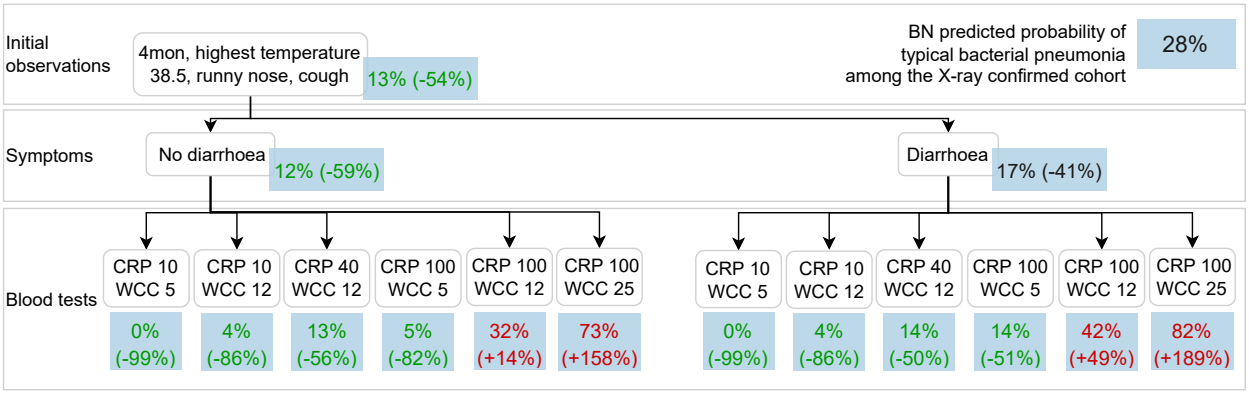

Fig 9. *Clinical scenario 1, BN v4.5.* This use case presents an infant with X-ray confirmed pneumonia, moderately high temperature, cough and runny nose. The BN predicts a low probability of typical bacterial pneumonia for this child, and a reported history of diarrhoea alone does not largely affect the prediction. The target prediction remains low unless both the CRP and WCC are moderately elevated.

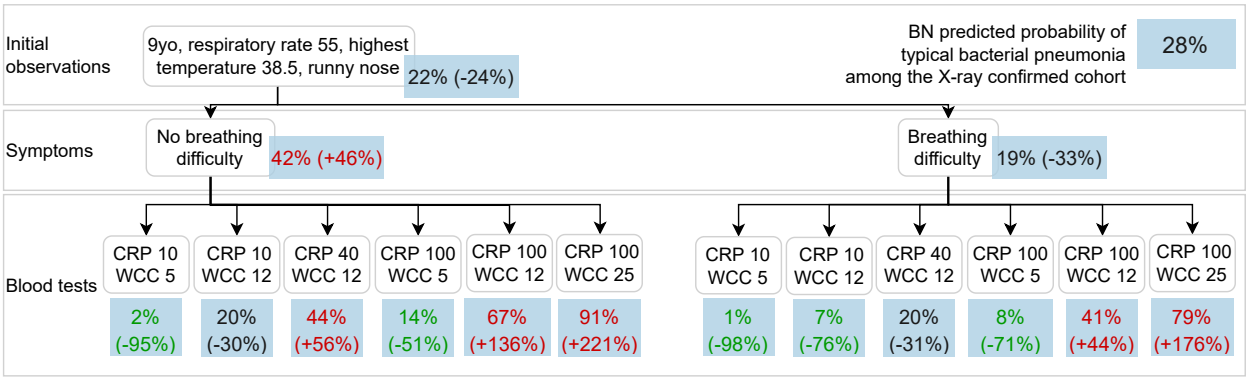

Fig 10. *Clinical scenario 2, BN v4.5.* This scenario presents a school-aged child with X-ray confirmed pneumonia and with moderately high temperature, runny nose and high respiratory rate. The BN predicts a 22% probability of typical bacterial pneumonia for this child, and this prediction increases to 42% if the child has no breathing difficulty. Moderate and very high CRP and WCC further increase the probability of a typical bacterial cause for this pneumonia.

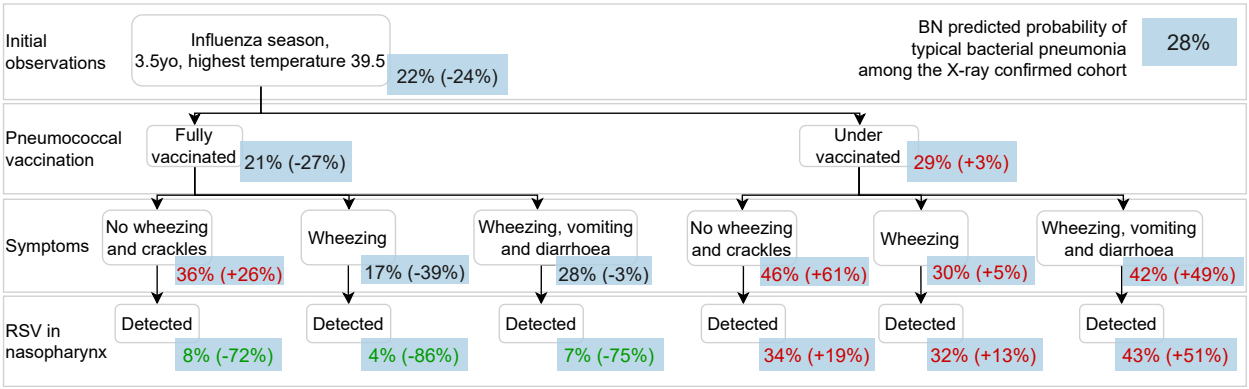

Fig 11. *Clinical scenario 3, BN v4.5.* This scenario shows a pre-school aged child with X-ray confirmed pneumonia and very high temperature during the influenza season. The probability of typical bacterial pneumonia increases if the child is under-vaccinated for pneumococcus, regardless of whether RSV is detected from their nasopharynx or not. If fully-vaccinated, the presence of wheezing will lead to a decrease in the predicted probability of typical bacterial pneumonia, and this prediction will drop to only 4% if RSV is detected in the child's nasopharynx.

4 Discussion

The diagnostic challenge of differentiating bacterial from non-bacterial pneumonia is the main driver of antibiotic use for treating pneumonia in children. We have built a causal BN that captures the complex epidemiological context of paediatric pneumonia and how this gives rise to the variable clinical manifestations observed in the hospital setting. Based on the combined use of domain expert knowledge and data, the model was designed to apply to a cohort of children with X-ray confirmed pneumonia who presented to a tertiary paediatric hospital in Australia. The resulting BN offers explainable and quantitative predictions on a range of variables of interest, including the diagnosis of bacterial pneumonia, detection of respiratory pathogens in the nasopharynx, and the clinical phenotype of a pneumonia episode. These outcomes can be utilized to derive recommendations to support more directed and judicious use of antimicrobials for this cohort.

Quantitatively, the model achieved stable and good performance (Table 1) in predicting the directly observable variables like the diagnosis of pleural effusion, positive blood culture and detection of RSV

and mycoplasma in the nasopharynx, based on 10-fold cross-validation. Qualitatively, we used commonly encountered scenarios to assess the potential usefulness of the model outputs in various clinical pictures. To our knowledge, this is the first causal model developed to help determine the causative pathogen for paediatric pneumonia. Our model framework and the methodological approach can be adapted beyond our context to broad respiratory infections and geographical and healthcare settings.

4.1 Using latent concepts to reveal the epidemiological context

When a child has X-ray confirmed pneumonia, the causative pathogen for pneumonia usually remains unobserved, i.e., latent. Patient background factors such as basic demographics, seasonality, vaccination and comorbidities can help predict the latent causative pathogen because they often causally influence the presence of a respiratory pathogen in the child's nasopharynx or respiratory tract which in turn predisposes to pneumonia. For example, RSV is more prevalent in influenza season owing to its elevated circulation in the community, and bacterial infection is more probable in a younger child due to the interplay between their developing immunity and respiratory microbiota [38, 39]. The difficulties faced by clinicians arises from the dynamic and complex interactions amongst these factors, as we saw in Fig. 11 for a young child who presents to hospital in the influenza season.

We addressed this challenge by firstly organising relevant variables under a causal framework (Fig. 4), and introducing several latent concepts such as the level of exposure and the susceptibility of a child to colonisation. This approach helped explain the potentially complex role of each variable in the problem domain. For example, older children may have lower susceptibility to bacterial colonisation as well as to pneumonia progression, with both leading to a reduced probability of typical bacterial pneumonia [39]. However, they are also more likely to have a high exposure through attendance at childcare or school [40], leading to an increase in the nasopharyngeal presence of all pathogens, thereby increasing the potential for subsequent infection. Being explicit about these underlying mechanisms allowed us to capture and understand the role of each factor in the context of each pneumonia episode, which in turn gave us further insight into the epidemiological context and more precise predictions of the causative pathogen for pneumonia.

4.2 Mapping clinical and laboratory evidence to explicit causal paths

Model predictions can be further improved by utilising additional information which is likely to be available shortly after presentation, namely, clinical observations and laboratory investigation results. By modelling the overlap and divergence in the pathophysiological pathways indicated by each piece of clinical and laboratory evidence (Fig. 6), the causal BN approach can estimate the episode-specific probability of each causative pathogen based on any input variables available. As illustrated in clinical scenario 1 (Fig. 9), while diarrhoea can be a feature of either viral or bacterial pneumonia, knowing about diarrhoea alone does not greatly affect the predicted probability of typical bacterial pneumonia. However, when additional information such as CRP and WCC becomes available, the prediction can vary from close to 0% up to 82%.

A challenge faced by clinicians is the variable predictive value of specific clinical and laboratory observations when the underlying epidemiology changes [14]. More specifically, typical bacterial pneumonia is more likely to drive elevated WCC and CRP [30], however, elevated biomarkers can also be caused by viral pathogens, especially when disease manifestation are severe, leading to reduced specificity for bacterial infection. Our BN addressed this issue by introducing a variable to summarise the clinical phenotype for each episode, which functioned to separate the influence of the causative pathogen from the clinical progression of the presenting episode based on the patient's susceptibility to progression and the duration and severity of the presenting illness. This was illustrated in Scenario 1 (Fig. 9) and Scenario 2 (Fig. 10), in which the same blood results could lead to significantly different predictions of typical bacterial pneumonia for the two demographically and clinically distinct presenting episodes.

4.3 The challenge and need to better understand pathogenicity

It became apparent since early in the model development (after Workshop 1) that the pathogenicity of different organisms plays a central role in driving the probability that a pathogen is causative. There are a large number of organisms (viral and bacterial) that are capable of causing pneumonia, while there are also

frequent detections of multiple potentially pathogenic organisms in the nasopharynx. Hence, it is important to understand the prevalence of relevant organisms in the respiratory tract (including the nasopharynx) as well as the probability of being causative of pneumonia given they are present there. These issues affect how the model should be structured (e.g., to what extent we should separate or group pathogens) and how we should parameterise the model (e.g., given the co-presence of more than one pathogen, which one is the more likely to be the causative agent). Many of these issues remain the subject of debate and ultimately unresolved by the scientific community [41].

We designed a series of survey questions to elicit the current understanding and knowledge of domain experts (Survey 2, 3 and 5). Finding the best way to frame questions proved challenging given the complexity of the problem domain. Consequently, we explored a variety of approaches based on modelling needs, facilitated by graphical illustration, data visualisation, and various question formats. This process also helped clarify the modelling requirements themselves, helping us to better manage our current imperfect understanding of the problem domain. In the areas of greatest need, sensitivity analyses were subsequently conducted to understand how much the uncertainty in our understanding could impact on the BN predicted outcomes (Fig. 8).

4.4 Implications for providing decision support for antibiotic prescription

Satisfactory numeric predictive performance has been achieved by the BN in predicting directly observable variables, especially those used to operationally define bacterial pneumonia. The model was qualitatively reviewed in various forms, including in Workshops 3 and 4 and Survey 4. The review outcomes from our experts from diverse domain areas provided confidence in the potential clinical value that this BN offers as a decision support tool.

Among the 230 children enrolled in our cohort of X-ray confirmed pneumonia, 99% received antibiotics during their hospital stay, although only 11% were confirmed with a clinical diagnosis of bacterial pneumonia [30]. This indicates that the current approach to clinically diagnose bacterial pneumonia fails to provide clinicians the information they need to confidently defer antibiotics. Applied to this cohort, the outcome BN predicted that 28% had typical bacterial pneumonia (either bacterial or viral-bacterial co-infection), indicating the potential to greatly rationalise use of antibiotics.

The three scenarios in § 3.4 provide an insight into how we envisage the model outputs could support decisions under various clinical contexts, leading to numerous questions about how model outputs should be presented. Even if causal models have good performance and produce explainable outputs, they will not be useful tools unless their outputs can be presented in a form that end-users can understand, interpret and use. The barriers to implementation and uptake that decision support tools including BNs face are widely reported [42, 43]

4.5 Study limitations and future direction

The BN needs further validation before it can be clinically implemented, particularly examining the performance of the model in additional cohorts that differ in both key contextual factors and in the availability and operational definitions of superficially equivalent variables. Examples of such differences include: whether the definition of pneumonia is distinct across studies; whether a different spectrum of respiratory infections are involved; and how the clinical epidemiology differs based on the geographically distinct transmission dynamics or variations across healthcare systems. The results of these additional cohort validations should in turn feed back and inform how the BN should be further developed and adapted to other settings.

It is important to note there is a distinction between the actual presence of an organism at a body site and its successful detection by sampling at that body site; we only ever know about the former via evidence of the latter. In the model, we used the detection of organisms as a surrogate for their actual presence in the nasopharynx, in order to reduce the number of latent variables. For our model development, this data was collected prospectively and PCR testing was performed systematically to identify relevant pathogens over and above the routine clinical practice, generating high quality information with little missing data. However, the quality of model predictions will not be as robust if data from routine clinical review and investigation is used alone. Future work should consider the sensitivity and specificity of the laboratory techniques used to capture this information.

The BN can be further improved using richer data with a higher level of granularity. For example, information on the severity of reported symptoms was limited. Furthermore, other comorbidities were identified but are yet to be included in the model, such as the history of neurological conditions that may affect susceptibility to colonisation and disease progression. There is also the potential to include further information on the use of supplemental fluid, intensive care management and patient progression over time, which would enable the model to predict the likely effects of specific interventions. Integrating such models into electronic medical record systems to enable efficient data capture is an important future direction, as it would significantly improve the ability to evaluate, maintain and use these models and, ultimately, to derive the most value from them.

Acknowledgments

This project was supported by Perth Children Hospital Foundation (Grant ID 9900). YW is supported by the Western Australian Health Translation Network Early Career Fellowship and the Australian Government's Medical Research Future Fund (MRFF) as part of the Rapid Applied Research Translation program. CCB is supported by an Investigator grant from the National Health and Medical Research Council (NHMRC, APP1173163). TS is supported by a NHMRC Career Development Fellowship (GNT1111657). LK is supported by a Perron Foundation fellowship and NHMRC Ideas Grant (APP2003876). AOM is supported by a NHMRC Postgraduate Scholarship (APP1191465) and an Australian Government Research Training Program Fees Offset. CM is supported by a RAINE Clinician Research Scholarship. PCMW is supported by a NHMRC Investigator grant (APP1197335).

References

1. McIntosh K. Community-Acquired Pneumonia in Children. New England Journal of Medicine. 2002;346(6):429–437. doi:10.1056/NEJMra011994.
2. Gupta GR. Tackling pneumonia and diarrhoea: the deadliest diseases for the world's poorest children. The Lancet. 2012;379(9832):2123–2124. doi:10.1016/S0140-6736(12)60907-6.
3. McCollum ED, King C, Hammitt LL, Ginsburg AS, Colbourn T, Baqui AH, et al. Reduction of childhood pneumonia mortality in the Sustainable Development era. The Lancet Respiratory medicine. 2016;4(12):932–933.
4. Goodman D, Crocker ME, Pervaiz F, McCollum ED, Steenland K, Simkovich SM, et al. Challenges in the diagnosis of paediatric pneumonia in intervention field trials: recommendations from a pneumonia field trial working group. The Lancet Respiratory Medicine. 2019;7(12):1068–1083. doi:https://doi.org/10.1016/S2213-2600(19)30249-8.
5. Bhuiyan MU, Snelling TL, West R, Lang J, Rahman T, Granland C, et al. The contribution of viruses and bacteria to community-acquired pneumonia in vaccinated children: a case–control study. Thorax. 2019;74(3):261–269.
6. O'Brien KL, Baggett HC, Brooks WA, Feikin DR, Hammitt LL, Higdon MM, et al. Causes of severe pneumonia requiring hospital admission in children without HIV infection from Africa and Asia: the PERCH multi-country case-control study. The Lancet. 2019;394(10200):757–779.
7. Cohen ML. Epidemiology of drug resistance: implications for a post—antimicrobial era. Science. 1992;257(5073):1050–1055.
8. Shah SN, Bachur RG, Simel DL, Neuman MI. Does this child have pneumonia?: the rational clinical examination systematic review. JAMA. 2017;318(5):462–471.
9. Brealey JC, Sly PD, Young PR, Chappell KJ. Viral bacterial co-infection of the respiratory tract during early childhood. FEMS microbiology letters. 2015;362(10):fnv062.
10. Bradley JS. Management of community-acquired pediatric pneumonia in an era of increasing antibiotic resistance and conjugate vaccines. The Pediatric infectious disease journal. 2002;21(6):592–598.

11. Lynch T, Platt R, Gouin S, Larson C, Patenaude Y. Can we predict which children with clinically suspected pneumonia will have the presence of focal infiltrates on chest radiographs? Pediatrics. 2004;113(3):e186–e189.

12. Mahabee-Gittens EM, Grupp-Phelan J, Brody AS, Donnelly LF, Bracey SEA, Duma EM, et al. Identifying children with pneumonia in the emergency department. Clinical Pediatrics. 2005;44(5):427–435.

13. Berg AS, Inchley CS, Fjaerli HO, Leegaard TM, Lindbaek M, Nakstad B. Clinical features and inflammatory markers in pediatric pneumonia: a prospective study. European journal of pediatrics. 2017;176(5):629–638.

14. Joffe MD, Alpern ER. Occult pneumococcal bacteremia: a review. Pediatric emergency care. 2010;26(6):448–454.

15. Michelow IC, Olsen K, Lozano J, Rollins NK, Duffy LB, Ziegler T, et al. Epidemiology and clinical characteristics of community-acquired pneumonia in hospitalized children. Pediatrics. 2004;113(4):701–707.

16. Harris M, Clark J, Coote N, Fletcher P, Harnden A, McKean M, et al. British Thoracic Society guidelines for the management of community acquired pneumonia in children: update 2011. Thorax. 2011;66(Suppl 2):ii1–ii23.

17. Pratt MT, Abdalla T, Richmond PC, Moore HC, Snelling TL, Blyth CC, et al. Prevalence of respiratory viruses in community-acquired pneumonia in children: a systematic review and meta-analysis. The Lancet Child & Adolescent Health. 2022;.

18. Mathur S, Fuchs A, Bielicki J, Van Den Anker J, Sharland M. Antibiotic use for community-acquired pneumonia in neonates and children: WHO evidence review. Paediatrics and international child health. 2018;38(sup1):S66–S75.

19. Clark JE. Determining the microbiological cause of a chest infection. Archives of Disease in Childhood. 2015;100(2):193–197.

20. Moreno L, Krishnan JA, Duran P, Ferrero F. Development and validation of a clinical prediction rule to distinguish bacterial from viral pneumonia in children. Pediatric pulmonology. 2006;41(4):331–337.

21. Van De Maat J, Nieboer D, Thompson M, Lakhanpaul M, Moll H, Oostenbrink R. Can clinical prediction models assess antibiotic need in childhood pneumonia? A validation study in paediatric emergency care. PloS one. 2019;14(6):e0217570.

22. Bilkis MD, Gorgal N, Carbone M, Vazquez M, Albanese P, Branda MC, et al. Validation and development of a clinical prediction rule in clinically suspected community-acquired pneumonia. Pediatric emergency care. 2010;26(6):399–405.

23. Knop M, Weber S, Mueller M, Niehaves B, et al. Human Factors and Technological Characteristics Influencing the Interaction of Medical Professionals With Artificial Intelligence–Enabled Clinical Decision Support Systems: Literature Review. JMIR Human Factors. 2022;9(1):e28639.

24. Kyrimi E, McLachlan S, Dube K, Neves MR, Fahmi A, Fenton N. A comprehensive scoping review of bayesian networks in healthcare: Past, present and future. Artificial Intelligence in Medicine. 2021;117:102108.

25. McLachlan S, Dube K, Hitman GA, Fenton NE, Kyrimi E. Bayesian networks in healthcare: Distribution by medical condition. Artificial intelligence in medicine. 2020;107:101912.

26. Pearl J. Embracing causality in default reasoning. Artificial Intelligence. 1988;35(2):259–271.

27. Korb KB, Nicholson AE. Bayesian Artificial Intelligence. 2nd ed. Boca Raton: Chapman & Hall / CRC Press; 2011.

28. Wu Y, McLeod C, Blyth C, Bowen A, Martin A, Nicholson A, et al. Predicting the causative pathogen among children with osteomyelitis using Bayesian networks–improving antibiotic selection in clinical practice. Artificial Intelligence in Medicine. 2020;107:101895.

29. Ramsay JA, Mascaro S, Campbell AJ, Foley D, Mace AO, Ingram P, et al. Urinary tract infections in children: building a causal model-based decision support tool for diagnosis with domain knowledge and prospective data. medRxiv. 2022;.

30. Bhuiyan MU, Blyth CC, West R, Lang J, Rahman T, Granland C, et al. Combination of clinical symptoms and blood biomarkers can improve discrimination between bacterial or viral community-acquired pneumonia in children. BMC pulmonary medicine. 2019;19(1):1–9.

31. Dempster AP, Laird NM, Rubin DB. Maximum likelihood from incomplete data via the EM algorithm. Journal of the Royal Statistical Society: Series B (Methodological). 1977;39(1):1–22.

32. Good IJ. Rational decisions. In: Breakthroughs in statistics. Springer; 1992. p. 365–377.

33. Sobol IM. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. Mathematics and computers in simulation. 2001;55(1-3):271–280.

34. Borgonovo E, et al. Sensitivity analysis. An Introduction for the Management Scientist International Series in Operations Research and Management Science Cham, Switzerland: Springer. 2017;.

35. Spirtes P, Glymour C, Scheines R. Causation, Prediction, and Search. vol. 81 of Lecture Notes in Statistics. Springer New York;. Available from: `http://link.springer.com/10.1007/978-1-4612-2748-9`.

36. R Core Team. R: A Language and Environment for Statistical Computing; 2021. Available from: `https://www.R-project.org/`.

37. Van Rossum G, Drake FL. Python 3 Reference Manual. Scotts Valley, CA: CreateSpace; 2009.

38. Bosch AA, Piters WAdS, van Houten MA, Chu MLJ, Biesbroek G, Kool J, et al. Maturation of the infant respiratory microbiota, environmental drivers, and health consequences. A prospective cohort study. American journal of respiratory and critical care medicine. 2017;196(12):1582–1590.

39. Ramos-Sevillano E, Ercoli G, Brown JS. Mechanisms of naturally acquired immunity to Streptococcus pneumoniae. Frontiers in immunology. 2019;10:358.

40. Alexandrino AS, Santos R, Melo C, Bastos JM. Risk factors for respiratory infections among children attending day care centres. Family practice. 2016;33(2):161–166.

41. Douros K, Kotzia D, Kottaridi C, Giotas A, Boutopoulou B, Bozas E, et al. Evidence for respiratory viruses interactions in asymptomatic preschool-aged children. Allergologia et Immunopathologia. 2019;47(3):260–264.

42. Moxey A, Robertson J, Newby D, Hains I, Williamson M, Pearson SA. Computerized clinical decision support for prescribing: provision does not guarantee uptake. Journal of the American Medical Informatics Association. 2010;17(1):25–33.

43. Kyrimi E, Dube K, Fenton N, Fahmi A, Neves MR, Marsh W, et al. Bayesian networks in healthcare: What is preventing their adoption? Artificial Intelligence in Medicine. 2021;116:102079.