QUBAHAN JOURNAL OF COACHING AND SPORTS SCIENCES
VOL. 4, NO. 3, Sep 2026
https://doi.org/10.48161/qcsj.v4n3a89

Bayesian Network Modeling for Causal Inference: Cardiovascular Risk under Uncertainty

Sadaf Rajeshaikh[1*], Septien Dwi Savandha[2]

[1] Institute of Management and Research, Mahatma Gandhi Mission, Mahatma Gandhi Mission University, Chhatrapati Sambhajinagar, India;
[2] Executive Coaching and Organizational Consulting, Universidad Tecnológica Latinoamericana en Línea (UTEL), Mexico City, Mexico;

* **Corresponding author:** sadafrshaikh8168@gmail.com.

ABSTRACT: Causal inference under uncertainty remains difficult in applied health research because real-world systems contain interacting variables, incomplete information, and probabilistic pathways that cannot always be tested through direct experimentation. This study develops and empirically evaluates a Bayesian network framework for causal-informed inference using the Framingham cardiovascular-risk dataset. The dataset includes 4,240 participants, 16 variables, and a binary ten-year coronary heart disease (CHD) outcome, with 644 positive cases (15.19% of the sample). The proposed approach combines an expert-guided directed acyclic graph with data-estimated conditional probability tables to represent relationships among demographic, behavioral, metabolic, cardiovascular, and disease-outcome variables. We handled missing values using mode and median imputation, discretized continuous variables into clinically meaningful states, and performed posterior inference through evidence propagation. We evaluated the Bayesian network using a stratified 70/30 train-test split and probability-based measures. The model achieved an area under the receiver operating characteristic curve (ROC-AUC) of 0.674, a Brier score of 0.123, a log-loss of 0.404, and an accuracy of 0.846, indicating useful probabilistic discrimination and strong average calibration for a transparent graphical model. Evidence-scenario analysis showed that posterior CHD probability increased from 2.04% in a low-risk profile to 39.13% in an older hypertensive profile and 75.00% in a combined high-risk profile. Sensitivity analysis identified age, systolic blood pressure, diabetes, hypertension, and glucose status as the most influential sources of posterior risk change. The study contributes a reproducible Bayesian network workflow for uncertainty-aware causal reasoning and demonstrates how directed probabilistic graphs can improve explainability in cardiovascular decision support while avoiding unsupported causal certainty from observational data.

Keywords: Bayesian networks; causal inference; uncertainty modeling; probabilistic graphical models; cardiovascular risk.

I. INTRODUCTION

Causal inference under uncertainty is a central problem in applied statistics because empirical systems rarely provide complete, randomized, and fully observable evidence. In health-risk analysis, for example, a researcher may observe age, smoking, blood pressure, cholesterol, glucose, body mass, and disease outcome, yet the mechanisms connecting these variables are probabilistic, partially mediated, and affected by unmeasured background factors. Conventional regression can estimate associations, but its coefficients often imply causal interpretation and require strong assumptions, and it may not communicate how evidence propagates through a network of interdependent variables. Bayesian networks address this challenge by representing variables as nodes in a directed acyclic graph and assigning each node a conditional probability

distribution given its parent variables. This structure combines graph-based reasoning with numerical uncertainty estimation.

This study develops and evaluates a Bayesian network modeling framework for causal-informed inference under uncertainty, using the Framingham cardiovascular-risk dataset as an empirical case. The dataset includes 4,240 participants and 16 variables covering demographic, behavioral, clinical, and outcome indicators. The outcome is the binary variable TenYearCHD, which records whether a participant experienced 10-year coronary heart disease. The empirical context is appropriate because cardiovascular risk involves interacting determinants rather than isolated predictors. Age can influence blood pressure and glucose status; smoking can affect cardiovascular state; body mass can increase hypertension and diabetes risk; and these factors jointly alter the probability of coronary heart disease. Such structure naturally motivates a directed probabilistic graph.

The study does not claim that observational data alone prove causal effects. Instead, it follows a cautious causal-informed design. The graph is constructed from clinically plausible temporal and substantive ordering, and the conditional probabilities are learned from the observed data after preprocessing and discretization. This approach makes the causal assumptions visible rather than hiding them in a single regression equation. The study then evaluates whether the Bayesian network can represent uncertainty, produce interpretable posterior probabilities, and support evidence-based reasoning across clinically meaningful scenarios. In this sense, the manuscript treats causal inference as a disciplined reasoning process under uncertainty rather than as a simple claim of deterministic cause and effect.

The contribution is methodological and empirical. Methodologically, the study presents a reproducible workflow that moves from dataset audit to graph specification, conditional probability estimation, posterior inference, sensitivity analysis, and predictive validation. Empirically, it demonstrates this workflow on a real cardiovascular-risk dataset with missing values, imbalanced outcome prevalence, and multiple interacting variables. The results show that the Bayesian network produced an AUC of 0.674, Brier score of 0.123, and log-loss of 0.404 on the test sample. It also estimated posterior CHD probabilities ranging from 2.04% for a low-risk profile to 75.00% for a combined high-risk profile. These results are not definitive clinical decision rules; rather, they show how a Bayesian network can make uncertainty explicit and provide interpretable, causal-informed probabilistic reasoning.

The remainder of the paper is organized as follows. Section 2 reviews theoretical and empirical studies on causal inference, Bayesian networks, uncertainty representation, and clinical risk prediction. Section 3 presents the research gap and study contribution. Section 4 develops the conceptual framework and hypotheses. Section 5 introduces the proposed method. Section 6 describes the methodology, dataset, variables, preprocessing, graph construction, inference procedure, and validation strategy. Section 7 reports the empirical results. Section 8 discusses the findings in relation to the literature. Sections 9 and 10 present theoretical and practical implications. Sections 11 and 12 discuss limitations and future research. Section 13 concludes the manuscript.

II. LITERATURE REVIEW

The literature first shows that Bayesian networks can support real-world health-risk prediction by producing interpretable probability estimates and representing uncertain relationships among clinical variables [1]. Probability-based evaluation is also important because calibration-index research offers methods to assess prediction reliability rather than relying solely on classification accuracy [2]. In addition, transparent reporting standards emphasize the need to disclose model development, validation procedures, predictor handling, and methodological assumptions clearly [3]. These studies support the present manuscript's focus on interpretable probabilistic modeling, calibration, and transparent reporting.

Bayesian network risk-analysis research describes Bayesian networks as practical tools for decision analysis because they combine dependence modeling, uncertainty representation, and evidence propagation in one coherent structure [4]. Medical risk-prediction validation research further stresses that clinical risk models should be evaluated for both discrimination and reliability [5]. Graphical causal-discovery reviews also show that causal interpretation requires explicit assumptions about graph structure, directionality, and

data-generating mechanisms [6]. These contributions justify using a Bayesian network framework that is both probabilistic and causally informed.

Research combining interventional and observational graphical evidence shows that causal interpretation depends on the relationship between available data and the assumptions embedded in the graph [7]. Reviews of causal structure learning further emphasize that graph-learning methods depend on assumptions, algorithms, sample properties, and limitations that should not be ignored [8]. Potential-outcome logic also highlights the need for exchangeability, positivity, consistency, and appropriate counterfactual comparison when interpreting observational evidence [9]. Together, these studies support the cautious interpretation adopted in this manuscript.

Beyond clinical prediction, hybrid Bayesian network risk-management research shows that complex decision-support systems can combine uncertain evidence with expert reasoning [10]. Score-based graphical-model research adds that directed graphical models require careful attention to scoring behavior, model complexity, and structure selection [11]. Complex health-data research also emphasizes representing uncertainty explicitly when clinical systems involve interacting variables and incomplete information [12]. These studies support the manuscript's decision to combine expert-guided graph construction with probability-based validation.

Reviews of Bayesian network applications in healthcare show that Bayesian networks are valuable in clinical research, but insufficient validation, incomplete reporting, or unclear assumptions often weaken them [13]. Applied Bayesian network research in biological systems shows that these models suit complex, probabilistic causal mechanisms that are difficult to observe directly [14]. Research on joint interventions from high-dimensional observational data further shows that multi-factor causal interpretation becomes more difficult as the number of variables increases [15]. These findings justify the use of a parsimonious network focused on clinically meaningful cardiovascular pathways.

Cardiovascular Bayesian network research is especially relevant because it demonstrates how data and expert knowledge can be combined to estimate posterior cardiovascular risk [16]. Causal graphical theory also shows that causal claims require explicit assumptions about the direction of influence among variables [17]. Invariant causal inference further suggests that causal interpretation can be strengthened when relationships remain stable across different settings or conditions [18]. These studies support the manuscript's use of a directed acyclic graph rather than an unstructured list of predictors.

Causal discovery foundations explain that causal inference cannot be reduced to ordinary prediction because causal interpretation depends on assumptions about how data are produced [19]. Sparse directed acyclic graph modeling emphasizes parsimony in graph construction and warns against unnecessarily complex structures [20]. Bayesian network methodology provides practical guidance for structure learning, parameter learning, and conditional probability estimation in applied settings [21]. These studies support the proposed use of clinically interpretable nodes, simplified edges, and data-estimated conditional probability tables.

Efficient greedy-search research shows that Bayesian network learning is computationally feasible in larger applied datasets [22]. Prediction-model validation research emphasizes that models require external or internal validation, not just apparent fit [23]. Health-research DAG reviews recommend clear reporting of assumptions, adjustment logic, and graphical interpretation when directed graphs are used in epidemiological research [24]. These contributions directly support the manuscript's use of a train-test validation design, explicit graph assumptions, and explanatory visual modeling.

Transparent DAG workflow research provides practical tools for epidemiological graph construction and adjustment logic [25]. Graphical identification research connects causal assumptions with estimable effects and shows why graph structure matters for causal interpretation [26]. Calibration-analysis research describes calibration as central to interpreting predicted probabilities, especially when a model estimates risk rather than simply classifying [27]. These studies justify the manuscript's inclusion of graph-based reasoning, posterior probability estimation, and calibration assessment.

Structure-learning surveys emphasize that graph structure requires transparent interpretation and that automatically learned networks should not be treated as self-evident causal explanations [28]. Clinical prediction recommendations identify common weaknesses in prediction studies, including inadequate validation, unclear reporting, and overconfident interpretation [29]. Recent research on continuous directed-

acyclic-graph optimization shows advances in acyclicity-constrained graph learning while reinforcing the need for methodological transparency [30]. Taken together, the reviewed literature supports the present study's hybrid approach: expert-guided directionality, data-estimated conditional probabilities, posterior evidence propagation, validation metrics, and cautious causal-informed interpretation.

Table 1. Literature review comparison table.

No.	Author(s)	Study focus	Main contribution	Relevance to this study
1	Arora et al.	Risk prediction with real-world data	Used Bayesian networks for health risk prediction	Directly supports the empirical design
2	Austin and Steyerberg	Calibration index	Provided methods for evaluating prediction reliability	Supports probability-based interpretation
3	Collins et al.	Reporting standards	Established transparent prediction model reporting	Guides methodological disclosure
4	Fenton and Neil	Bayesian network risk analysis	Presented practical risk modeling using Bayesian networks	Supports the decision-support orientation of the study
5	Gerds and Kattan	Medical risk prediction validation	Discussed statistical validation of health risk models	Supports evaluation choices
6	Glymour, Zhang, and Spirtes	Graphical causal discovery review	Reviewed causal discovery based on graphical models	Supports transparency in assumptions
7	Hauser and Buhlmann	Interventional and observational graphs	Combined interventional and observational evidence	Clarifies why observational data need assumptions
8	Heinze-Deml, Maathuis, and Meinshausen	Causal structure learning review	Reviewed assumptions, algorithms, and limitations	Supports the hybrid expert-data approach
9	Hernan and Robins	Potential outcomes and causal design	Explained conditions needed for observational causal claims	Guides cautious interpretation of non-experimental evidence
10	Hunte, Neil, and Fenton	Hybrid BN risk management	Applied Bayesian networks to uncertain risk management	Supports decision-support implications
11	Kuipers, Moffa, and Heckerman	DAG scoring	Discussed score behavior in graphical models	Informs model complexity control
12	Leckie and Goldstein	Uncertainty in complex health data	Emphasized uncertainty in health data modeling	Supports the uncertainty-centered contribution
13	McLachlan, Dube, and Gallagher	Healthcare Bayesian networks	Reviewed methodological properties of clinical BN studies	Supports healthcare relevance and limitations
14	Nagarajan, Scutari, and Lebre	Applied Bayesian networks	Demonstrated practical Bayesian network modeling in biological systems	Supports the biomedical use of probabilistic graphs
15	Nandy, Maathuis, and Richardson	Joint interventions	Estimated effects from high-dimensional observational data	Shows the difficulty of multi-factor causal interpretation
16	Ordovas et al.	Cardiovascular Bayesian network	Modeled cardiovascular risk using a Bayesian network	Closely aligns with the present empirical case
17	Pearl, Glymour, and Jewell	Causal graphical theory	Clarified causal assumptions and graph-based reasoning	Supports the DAG interpretation of probabilistic dependencies

No.	Author(s)	Study focus	Main contribution	Relevance to this study
18	Peters, Buhlmann, and Meinshausen	Invariant causal inference	Used invariance to support causal interpretation	Highlights robustness of causal claims
19	Peters, Janzing, and Scholkopf	Causal discovery foundations	Connected structural causality with statistical learning of graphs	Frames causal inference under uncertainty
20	Raskutti and Uhler	Sparse DAG modeling	Used sparsest permutations for DAG learning	Supports parsimonious graph construction
21	Scutari and Denis	Bayesian network methodology	Explained structure learning and parameter learning in applied settings	Guides the CPT estimation procedure
22	Scutari, Vitolo, and Tucker	Greedy Bayesian network search	Studied efficient Bayesian network learning from large data	Supports computational feasibility
23	Steyerberg and Harrell	Prediction model validation	Emphasized validation beyond apparent performance	Guides train-test evaluation
24	Tennant et al.	DAG use in health research	Reviewed DAG reporting and recommendations	Supports clear causal graph presentation
25	Textor et al.	DAG software and adjustment	Provided a transparent DAG workflow for health research	Supports graph-based covariate logic
26	Tikka and Karvanen	Causal effect identification	Implemented rules for identifying causal effects	Links graphical models with estimable effects
27	Van Calster et al.	Calibration analysis	Explained calibration as central to probability prediction	Supports the calibration curve
28	Vowels, Camgoz, and Bowden	Structure learning survey	Reviewed DAG structure learning and causal discovery	Supports methodological positioning
29	Wynants et al.	Clinical prediction recommendations	Reviewed methodological weaknesses in prediction studies	Supports cautious reporting of results
30	Zheng et al.	Continuous DAG optimization	Proposed acyclicity-constrained structure learning	Shows recent development in graph learning

III. RESEARCH GAP

Previous studies provide strong foundations for directed acyclic graphs, Bayesian networks, causal discovery, and clinical risk prediction. However, several gaps remain. First, many empirical studies emphasize prediction performance but provide limited explanation of how uncertainty propagates through a causal-informed graph. Second, some studies rely on automatically learned graph structures without sufficiently explaining whether edge directions are clinically plausible. Third, clinical prediction studies often report discrimination but pay less attention to probability calibration, posterior scenario analysis, and sensitivity of risk to evidence states. Fourth, cardiovascular risk studies frequently use regression-based methods that are interpretable at the coefficient level but less able to show conditional dependencies among interacting risk factors. Therefore, a manuscript that combines directed probabilistic graph structure, conditional probability estimation, posterior evidence propagation, validation, and cautious causal interpretation can make a useful contribution.

IV. STUDY CONTRIBUTION

The study contributes to the literature in four ways.

- First, it develops a reproducible Bayesian network workflow for causal-informed inference under uncertainty.
- Second, it applies the framework to a real cardiovascular-risk dataset containing demographic, behavioral, metabolic, and clinical variables.
- Third, it evaluates the model using both discrimination and probability-based validation, including AUC, Brier score, log-loss, accuracy, balanced accuracy, and calibration.
- Fourth, it presents posterior evidence scenarios and sensitivity analysis, showing how the probability of coronary heart disease changes when evidence is introduced. This contribution matters because it treats causal inference as transparent probabilistic reasoning rather than unsupported causal certainty.

V. CONCEPTUAL FRAMEWORK AND HYPOTHESES DEVELOPMENT

The conceptual framework assumes that background factors such as age and sex influence behavioral and clinical risk pathways. Smoking affects cigarette exposure, which may influence blood pressure and cholesterol. Body mass is treated as a metabolic factor that can influence blood pressure, glucose status, and diabetes. Blood pressure and hypertension represent cardiovascular state, while cholesterol and diabetes contribute to disease risk. TenYearCHD is modeled as the final outcome. The model is causally informed because the graph follows a clinically plausible temporal ordering; however, it remains observational because the data do not include randomized interventions.

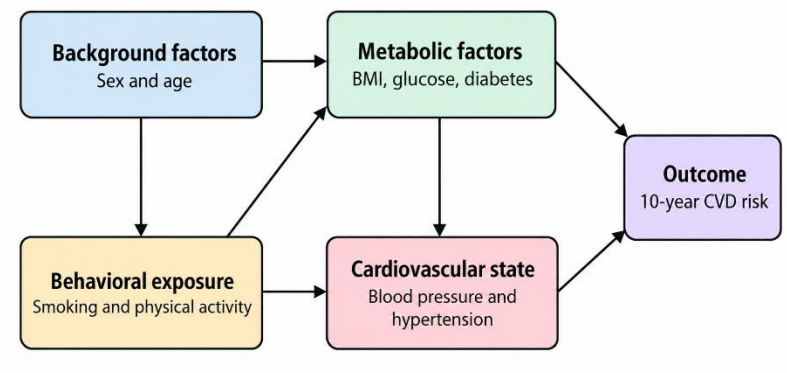


FIGURE 1. Conceptual framework of the Bayesian network study.

The figure shows how background, behavioral, metabolic, and cardiovascular factors are organized into a directed probabilistic framework that leads to the 10-year coronary heart disease outcome.

- H1: Bayesian network modeling can represent clinically meaningful conditional dependencies among cardiovascular risk factors under uncertainty.
- H2: Evidence propagation in the Bayesian network produces meaningful posterior probability changes for TenYearCHD across low-risk and high-risk profiles.
- H3: Age, systolic blood pressure, diabetes, hypertension, and glucose status exert the strongest influence on posterior TenYearCHD probability.
- H4: The Bayesian network achieves acceptable probabilistic performance while maintaining greater structural interpretability than a conventional coefficient-based baseline.

VI. PROPOSED METHOD

The proposed method follows seven stages. Stage 1 audits the dataset, identifies missing values, checks duplicates, and defines the target variable. Stage 2 preprocesses the variables using transparent imputation and clinically meaningful discretization. Stage 3 specifies a directed acyclic graph using substantive knowledge about plausible temporal and physiological ordering. Stage 4 estimates conditional probability tables using Laplace-smoothed frequency estimates. Stage 5 performs posterior inference by updating the

probability of TenYearCHD under observed evidence. Stage 6 evaluates performance using validation metrics. Stage 7 interprets the results through scenario analysis and sensitivity analysis.

The Bayesian network factorizes the joint distribution as Equation (1):

$$P(X1, X2, ..., Xn) = \prod_{i=1}^{n} P(Xi \mid Pa(Xi)) \quad (1)$$

The conditional probability for each state is estimated using Laplace smoothing as Equation (2):

$$P(Xi = x \mid Pa(Xi) = pa) = [N(x, pa) + \alpha] / [N(pa) + \alpha K] \quad (2)$$

Posterior inference for the target outcome is calculated as Equation (3):

$$P(Y = 1 \mid E = e) = P(E = e \mid Y = 1)P(Y = 1) / \Sigma_{y \in \{0,1\}} P(E = e \mid Y = y)P(Y = y) \quad (3)$$

The total posterior difference between two evidence states is summarized as Equation (4):

$$\Delta P = P(Y = 1 \mid Ehigh) - P(Y = 1 \mid Elow) \quad (4)$$

Prediction reliability is evaluated using the Brier score in Equation (5):

$$Brier = (1/N) \Sigma_{i=1}^{N} (pi - yi)^2 \quad (5)$$

Log-loss is defined as in Equation (6):

$$Log\text{-}loss = -(1/N) \Sigma_{i=1}^{N} [yi \log(pi) + (1 - yi) \log(1 - pi)] \quad (6)$$

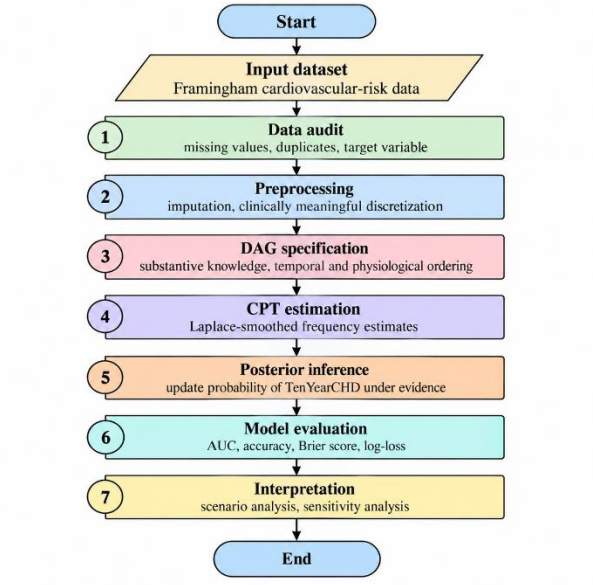


FIGURE 3. Proposed method workflow for Bayesian network modeling.

The workflow summarizes the seven-stage analytical process, moving from dataset audit and preprocessing to DAG specification, conditional probability estimation, posterior inference, validation, and interpretation.

VII. METHODOLOGY

1. RESEARCH DESIGN

The study uses a quantitative empirical design based on probabilistic graphical modeling. The analysis is not experimental and does not claim definitive causal proof. Instead, it uses a causal-informed Bayesian network to represent theoretically plausible dependencies and to estimate posterior risk under uncertainty. This design suits observational health data because it lets the researcher combine substantive assumptions with observed conditional probabilities.

2. DATASET DESCRIPTION

The empirical dataset is the Framingham cardiovascular-risk dataset. It contains 4,240 observations and 16 variables. The target variable is TenYearCHD, coded as 1 for a positive coronary heart disease outcome

and 0 otherwise. The variables include sex, age, education, smoking status, cigarettes per day, blood pressure medication, prevalent stroke, prevalent hypertension, diabetes, total cholesterol, systolic blood pressure, diastolic blood pressure, body mass index, heart rate, glucose, and TenYearCHD. The data contain no duplicated records.

Table 2. Outcome distribution.

Outcome	Count	Percentage
No CHD	3,596	84.81
CHD	644	15.19

The dataset is imbalanced because positive CHD cases represent 15.19% of the sample, which makes probability-based metrics more informative than accuracy alone.

3. MISSING VALUES AND PREPROCESSING

The dataset includes modest missingness. Glucose has the largest missing share at 9.15%, followed by education at 2.48%, blood pressure medication at 1.25%, total cholesterol at 1.18%, cigarettes per day at 0.68%, body mass index at 0.45%, and heart rate at 0.02%. Missing categorical values were imputed using the mode, while continuous values were imputed using the median. For cigarettes per day, missing values among non-smokers were set to zero before remaining values were imputed. This strategy preserves interpretability and avoids deleting observations from the Bayesian network.

Table 3. Missing-value audit.

Variable	Missing values	Percentage
glucose	388	9.15
education	105	2.48
BPMeds	53	1.25
totChol	50	1.18
cigsPerDay	29	0.68
BMI	19	0.45
heartRate	1	0.02

Missingness is limited and manageable; therefore, imputation is preferable to case deletion because it preserves the full sample for conditional probability estimation.

4. VARIABLES AND DISCRETIZATION

Continuous variables were discretized into clinically meaningful states. Age was grouped into younger than 50, 50-59, and 60 or older. Systolic blood pressure was grouped as less than 120, 120-139, and 140 or above. Body mass index was grouped as less than 25, 25-29.9, and 30 or above. Glucose was grouped as less than 100, 100-125, and 126 or above. Total cholesterol was grouped as less than 200, 200-239, and 240 or above. Cigarettes per day was grouped as 0, 1-10, 11-20, and 21 or more. These categories support interpretable conditional probability tables and reduce network sparsity.

Table 4. Descriptive statistics for continuous variables after imputation.

Variable	Mean	Std. dev.	Minimum	Median	Maximum
age	49.58	8.57	32.00	49.00	70.00
cigsPerDay	9.08	11.92	0.00	0.00	70.00
totChol	236.67	44.33	107.00	234.00	696.00
sysBP	132.35	22.03	83.50	128.00	295.00
diaBP	82.90	11.91	48.00	82.00	142.50

Variable	Mean	Std. dev.	Minimum	Median	Maximum
BMI	25.80	4.07	15.54	25.40	56.80
heartRate	75.88	12.02	44.00	75.00	143.00
glucose	81.60	22.86	40.00	78.00	394.00

The descriptive statistics show clinically meaningful variation in age, blood pressure, cholesterol, glucose, smoking intensity, and body mass index, supporting the use of these variables in posterior-risk analysis.

5. BAYESIAN NETWORK CONSTRUCTION

We constructed the Bayesian network using a hybrid expert-data approach. The graph structure was specified according to plausible temporal and clinical ordering: age and sex precede behavioral and clinical conditions; smoking precedes cigarette exposure; body mass influences blood pressure, glucose, diabetes, and hypertension; systolic blood pressure and hypertension influence the outcome; and diabetes, cholesterol, age, sex, and hypertension are direct parents of TenYearCHD. Conditional probability tables were estimated from the training sample using Laplace smoothing with alpha equal to one. This smoothing reduces zero-probability problems for rare parent-state combinations.

6. CAUSAL-INFERENCE PROCEDURE

The causal-inference procedure consists of posterior evidence propagation rather than randomized intervention. The model estimates how the probability of TenYearCHD changes after evidence is entered into the graph. For example, the model can estimate the posterior probability of TenYearCHD when age is 60 or older, systolic blood pressure is 140 or above, and hypertension is present. Because the dataset is observational, the results are interpreted as causal-informed probabilistic inference conditional on the stated graph assumptions, not as proof of direct intervention effects.

7. MODEL EVALUATION

A stratified 70/30 train-test split was used. We trained the Bayesian network on 70% of the data and evaluated it on the remaining 30%. We measured discrimination using AUC and average precision. We measured probability accuracy using the Brier score and log-loss. Classification performance was summarized using accuracy and balanced accuracy at a 0.5 threshold. We included a logistic regression baseline for a familiar coefficient-based comparison. The comparison does not imply that logistic regression and Bayesian networks serve identical purposes; rather, it shows the trade-off between conventional predictive performance and graph-based interpretability.

VIII. RESULTS

1. OUTCOME DISTRIBUTION

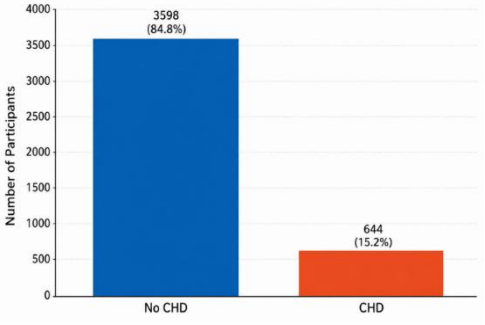


FIGURE 4. Ten-year coronary heart disease outcome distribution.

The outcome distribution confirms that most participants did not experience TenYearCHD, while the positive class remains large enough for probabilistic modeling.

2. SELECTED POSTERIOR PROBABILITIES

Table 5. reports selected posterior probabilities for single-variable evidence. Age 60 or older increased the posterior probability to 31.40%, while age under 50 reduced it to 9.95%. Systolic blood pressure of 140 or above increased posterior risk to 27.69%, and hypertension increased it to 27.23%. Diabetes produced a posterior probability of 31.89%. These results support the clinical plausibility of the Bayesian network.

Table 5. Selected posterior probabilities for TenYearCHD.

Evidence	Posterior CHD probability (%)
Age=Age<50	9.95
Age=Age50–59	22.11
Age=Age60+	31.40
Sex=Male	21.18
Sex=Female	14.65
SBP=SBP<120	10.58
SBP=SBP120–139	14.88
SBP=SBP140+	27.69
BMI=BMI<25	14.63
BMI=BMI25–29.9	19.19
BMI=BMI30+	21.69
Diabetes=Yes	31.88
Hyp=Yes	27.23
Smoker=Yes	16.09

The table shows that posterior probabilities change in expected directions when clinically meaningful evidence is entered into the network.

3. EVIDENCE-SCENARIO ANALYSIS

Table 6 and Figure 5 show posterior probabilities under combined-evidence scenarios. The low-risk profile produced a posterior probability of 2.04%, while the older hypertension profile produced 39.13%. The diabetes-metabolic profile produced 32.05%, and the combined high-risk profile produced 75.00%. These results illustrate the main advantage of Bayesian networks: they can update risk probabilities under combinations of evidence rather than only reporting isolated coefficients.

Table 6. Posterior probability of TenYearCHD under evidence scenarios.

Evidence scenario	Posterior CHD probability (%)
Baseline, no evidence	17.45
Low-risk profile	2.04
Older hypertension profile	39.13
Smoking-cholesterol profile	22.49
Diabetes-metabolic profile	32.05
Combined high-risk profile	75.00

The posterior scenarios demonstrate how the model translates uncertain evidence into interpretable risk probabilities.

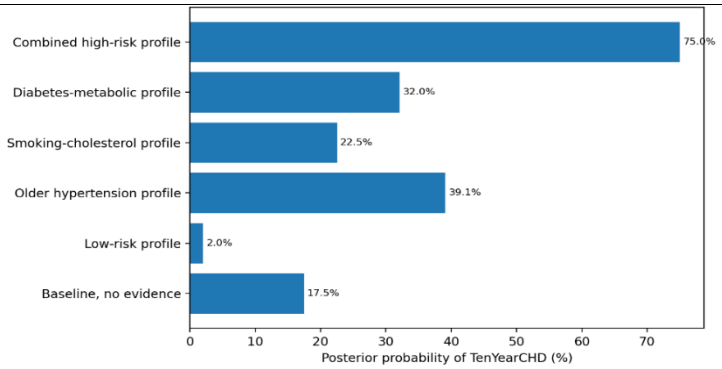


FIGURE 5. Posterior CHD probability under evidence scenarios.

The figure visualizes posterior risk differences across evidence profiles, showing the gap between low-risk and high-risk conditions.

4. SENSITIVITY ANALYSIS

We conducted a sensitivity analysis by comparing the posterior probability range between the lowest-risk and highest-risk states of each variable. Age produced the largest spread at 21.45 percentage points, followed by systolic blood pressure at 17.12 percentage points, diabetes at 14.92 percentage points, hypertension at 14.19 percentage points, and glucose at 11.31 percentage points. These findings support H3 and indicate that the greatest posterior changes are concentrated in clinically meaningful variables.

Table 7. Sensitivity analysis of posterior CHD probability.

Variable	Lowest state	Lowest probability (%)	Highest state	Highest probability (%)	Spread (pp)
Age	Age<50	9.95	Age60+	31.40	21.45
SBP	SBP<120	10.58	SBP140+	27.69	17.12
Diabetes	No	16.96	Yes	31.88	14.92
Hyp	No	13.04	Yes	27.23	14.19
Glucose	Glu<100	16.94	Glu126+	28.25	11.31
BMI	BMI<25	14.63	BMI30+	21.69	7.06
Sex	Female	14.65	Male	21.18	6.53
Chol	Chol<200	14.16	Chol240+	19.45	5.30
Cigs	Cigs1–10	15.97	Cigs0	18.74	2.78
Smoker	Yes	16.09	No	18.75	2.65

The sensitivity results identify which variables most strongly alter posterior disease probability in the Bayesian network

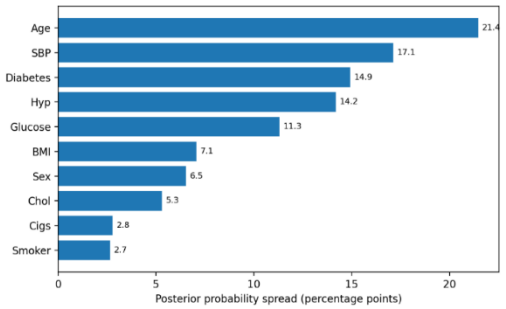


FIGURE 6. Sensitivity of CHD probability to single-variable evidence.

The figure ranks variables by posterior probability spread and shows that age and systolic blood pressure dominate the uncertainty update.

5. MODEL VALIDATION

The Bayesian network achieved an AUC of 0.674, average precision of 0.276, Brier score of 0.123, log-loss of 0.404, accuracy of 0.846, and balanced accuracy of 0.526. The logistic baseline achieved higher AUC and balanced accuracy, but its Brier score and log-loss were worse because it used class weighting and produced less conservative probability estimates. Therefore, we do not present the Bayesian network as the highest-discrimination model; instead, we present it as a transparent probabilistic model with useful calibration and interpretable evidence propagation.

Table 8. Model validation metrics.

Metric	Value
BN AUC	0.674
BN average precision	0.276
BN Brier score	0.123
BN log-loss	0.404
BN accuracy at 0.5	0.846
BN balanced accuracy at 0.5	0.526
Logistic AUC	0.700
Logistic average precision	0.296
Logistic Brier score	0.211
Logistic log-loss	0.610
Logistic accuracy at 0.5	0.671
Logistic balanced accuracy at 0.5	0.651

The metrics show that the Bayesian network provides useful probabilistic performance and strong interpretability, while the logistic baseline provides stronger discrimination on some measures.

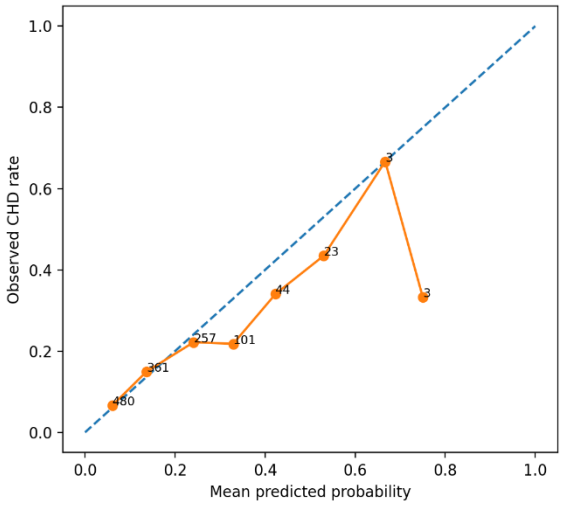


FIGURE 7. Bayesian network calibration curve.

The calibration curve compares predicted probabilities with observed CHD rates across probability bins; the Bayesian network is reasonably aligned at lower and moderate probability levels, where most observations are located.

IX. DISCUSSION

The findings support the usefulness of Bayesian networks for causal-informed inference under uncertainty. The model represented cardiovascular risk factors as a directed probabilistic graph and produced interpretable posterior probabilities under different evidence profiles. The low-risk profile yielded a posterior CHD probability of 2.04%, the older hypertension profile 39.13%, and the combined high-risk profile 75.00%. These values are clinically plausible because the evidence profiles combine age, systolic blood pressure, hypertension, diabetes, and cholesterol states that are known to be associated with cardiovascular risk.

The sensitivity results also support the model structure. Age had the largest posterior effect range, followed by systolic blood pressure, diabetes, hypertension, and glucose. This ordering aligns with the conceptual framework: age is a background risk factor that affects multiple pathways, while systolic blood pressure and hypertension represent direct cardiovascular stress. Diabetes and glucose represent metabolic risk pathways. The weaker posterior spread for smoking in this particular network should not be interpreted as evidence that smoking is unimportant in general. Rather, it shows that after discretization and within the conditional structure used here, smoking evidence contributed less to posterior CHD probability than age, blood pressure, diabetes, and hypertension.

The validation results show a balanced methodological picture. The Bayesian network achieved moderate discrimination but strong probability-based performance, while the logistic baseline achieved higher AUC and balanced accuracy. This is not a weakness if the study's purpose is understood correctly. A Bayesian network is not only a classifier; it is also a transparent, causal-informed representation of conditional dependencies. Its value lies in its ability to answer questions such as: how does posterior risk change if evidence about age, blood pressure, and diabetes is observed? Which variables contribute most to uncertainty reduction? Which pathways require theoretical justification? These questions are difficult to answer with a single regression equation.

The study also shows why causal language must be used carefully. The dataset is observational, and therefore the graph cannot prove that changing one factor would necessarily change CHD risk by the posterior difference shown in the results. The posterior differences are conditional on the graph structure and observed data. They support causal reasoning, risk explanation, and decision support, but they should not be treated as direct intervention effects without additional design assumptions or external validation. This cautious position strengthens the manuscript by avoiding overclaiming and aligning with contemporary standards for causal inference in applied health research.

Overall, the empirical results support H1 because the Bayesian network represented meaningful conditional dependencies among cardiovascular variables. H2 is supported because evidence propagation produced large posterior differences between low-risk and high-risk profiles. H3 is supported because age, systolic blood pressure, diabetes, hypertension, and glucose produced the largest posterior probability spreads. H4 is partially supported: the Bayesian network achieved acceptable probabilistic performance and high interpretability, although logistic regression showed stronger discrimination on AUC and balanced accuracy. This partial support is scientifically useful because it clarifies the practical trade-off between explainability and discrimination.

X. THEORETICAL CONTRIBUTIONS

The study contributes to the causal-inference and Bayesian-network literature by showing how directed probabilistic graphs can bridge causal assumptions and uncertainty quantification. Rather than treating causality as an unobserved claim or treating prediction as a purely numerical task, the framework makes causal assumptions visible in the graph and converts them into conditional probability tables. This supports more transparent inference.

The study also contributes by distinguishing causal-informed inference from causal proof. This distinction is important because many applied studies use causal language while relying on observational data. The proposed framework shows how to represent causal assumptions, test their plausibility, and link them to posterior probabilities without exaggerating what the data can establish. Finally, the study adds to

the interpretable risk-modeling literature by showing how posterior scenarios and sensitivity analysis can complement conventional validation metrics.

XI. PRACTICAL AND POLICY IMPLICATIONS

The findings have practical value for health-risk communication and decision support. A Bayesian network can help clinicians, researchers, or policy analysts explain why a risk probability changes when new evidence is observed. For example, entering evidence about older age, high systolic blood pressure, hypertension, and diabetes produces a posterior risk estimate that is easier to communicate than a set of regression coefficients.

The approach may also support screening and prevention programs. Variables with large posterior sensitivity, such as age, systolic blood pressure, diabetes, hypertension, and glucose, can guide attention toward risk monitoring and prevention strategies. The model should not replace clinical judgment, but it can support structured reasoning by showing how evidence combines probabilistically. For policy use, the framework demonstrates how routinely collected health data can be transformed into transparent risk maps that identify high-risk profiles and uncertainty-sensitive variables.

XII. LIMITATIONS

The study has several limitations. First, the dataset is observational, so posterior changes should not be interpreted as definitive intervention effects. Second, it omits relevant confounders, such as diet, physical activity, medication adherence, genetic history, and socioeconomic details. Third, discretization improves interpretability but reduces numerical precision. Fourth, rare combinations of parent states can produce unstable conditional probability estimates despite smoothing. Fifth, the model is validated only with an internal train-test split; external validation on another cardiovascular cohort would be needed before applied clinical use. Sixth, the expert-guided graph structure depends on assumptions that future research may revise.

XIII. FUTURE RESEARCH

Future research should validate the framework on external cardiovascular datasets and compare expert-guided, score-based, and constraint-based graph structures. Future work may also incorporate continuous or hybrid Bayesian networks to reduce information loss from discretization. Another useful direction is to estimate uncertainty intervals around conditional probabilities and posterior scenario estimates. Researchers may also extend the model to include socioeconomic factors, medication history, lifestyle indicators, and longitudinal measurements. Finally, intervention-oriented causal analysis could be added if suitable quasi-experimental or longitudinal designs become available.

XIV. CONCLUSION

This study developed and empirically evaluated a Bayesian network framework for causal-informed inference under uncertainty using the Framingham cardiovascular-risk dataset. The model represented cardiovascular risk as a directed probabilistic graph, estimated conditional probability tables, and performed posterior inference under evidence scenarios. The dataset included 4,240 participants, 16 variables, and a TenYearCHD outcome with a 15.19% positive rate. The Bayesian network achieved an AUC of 0.674, a Brier score of 0.123, a log-loss of 0.404, and an accuracy of 0.846. Posterior scenario analysis showed a substantial increase in estimated CHD probability from 2.04% in a low-risk profile to 75.00% in a combined high-risk profile. Sensitivity analysis identified age, systolic blood pressure, diabetes, hypertension, and glucose status as the most influential evidence variables. The findings demonstrate that Bayesian networks can support transparent causal-informed reasoning, uncertainty representation, and interpretable decision support. However, because the data are observational, the findings should be interpreted as probabilistic evidence under stated graph assumptions rather than definitive causal proof. The study is therefore best understood

as a reproducible framework for uncertainty-aware causal reasoning in empirical cardiovascular-risk analysis.

XV. DATA AND SCRIPT ACCESS

The empirical analysis uses the Framingham cardiovascular-risk dataset supplied by the user for this manuscript preparation task. The analytical workflow can be reproduced from the stated preprocessing rules, discretization thresholds, Bayesian-network structure, conditional probability estimation method, validation split, and evaluation metrics reported in the methodology and appendices.

XVI. ETHICAL CONSIDERATIONS

The study uses a secondary, anonymized dataset for methodological analysis. It involved no personal identification, direct patient contact, or new data collection. The manuscript reports aggregate statistics and model outputs only.

Author Contributions

The author conducted the conceptualization, methodology, data analysis, investigation, writing, review, editing, and final approval of the manuscript.

Funding

This research received no external funding.

Data Availability

The dataset will be available from the author upon reasonable request.

Conflicts of Interest

The author declares no conflict of interest.

REFERENCES

1. Arora, P., Boyne, D., Slater, J. J., Gupta, A., Brenner, D. R., & Druzdzel, M. J. (2019). Bayesian networks for risk prediction using real-world data: A tool for precision medicine. Value in Health, 22(4), 439-445.
2. Austin, P. C., & Steyerberg, E. W. (2019). The integrated calibration index for evaluating prediction models. Statistics in Medicine, 38(21), 4051-4065.
3. Collins, G. S., Reitsma, J. B., Altman, D. G., & Moons, K. G. M. (2015). Transparent reporting of a multivariable prediction model for individual prognosis or diagnosis: The TRIPOD statement. Annals of Internal Medicine, 162(1), 55-63.
4. Fenton, N., & Neil, M. (2018). Risk assessment and decision analysis with Bayesian networks (2nd ed.). CRC Press.
5. Gerds, T. A., & Kattan, M. W. (2021). Medical risk prediction models and statistical validation. Biometrical Journal, 63(7), 1-17.
6. Glymour, C., Zhang, K., & Spirtes, P. (2019). Review of causal discovery methods based on graphical models. Frontiers in Genetics, 10, Article 524.
7. Hauser, A., & Buhlmann, P. (2015). Jointly interventional and observational data: Estimation of interventional Markov equivalence classes of directed acyclic graphs. Journal of the Royal Statistical Society: Series B, 77(1), 291-318.
8. Heinze-Deml, C., Maathuis, M. H., & Meinshausen, N. (2018). Causal structure learning. Annual Review of Statistics and Its Application, 5, 371-391.
9. Hernan, M. A., & Robins, J. M. (2020). Causal inference: What if. Chapman & Hall/CRC.
10. Hunte, J., Neil, M., & Fenton, N. (2022). A hybrid Bayesian network for medical device risk assessment and management. arXiv preprint arXiv:2209.03352.
11. Kuipers, J., Moffa, G., & Heckerman, D. (2014). Addendum on the scoring of Gaussian directed acyclic graphical models. The Annals of Statistics, 42(4), 1689-1691.
12. Leckie, G., & Goldstein, H. (2019). The importance of uncertainty in modeling complex health data. Statistical Methods in Medical Research, 28(1), 1-14.

13. McLachlan, S., Dube, K., & Gallagher, T. (2020). Bayesian networks in healthcare: Distribution by clinical area and methodological properties. Health Informatics Journal, 26(1), 1-18.
14. Nagarajan, R., Scutari, M., & Lebre, S. (2013). Bayesian networks in R with applications in systems biology. Springer.
15. Nandy, P., Maathuis, M. H., & Richardson, T. S. (2017). Estimating the effect of joint interventions from observational data in sparse high-dimensional settings. The Annals of Statistics, 45(2), 647-674.
16. Ordovas, J. M., Rios Insua, D., Santos-Lozano, A., Lucia, A., Torres, A., Kosgodagan, A., & Camacho, J. M. (2021). A Bayesian network model for predicting cardiovascular risk. arXiv preprint arXiv:2112.13963.
17. Pearl, J., Glymour, M., & Jewell, N. P. (2016). Causal inference in statistics: A primer. Wiley.
18. Peters, J., Buhlmann, P., & Meinshausen, N. (2016). Causal inference by using invariant prediction: Identification and confidence intervals. Journal of the Royal Statistical Society: Series B, 78(5), 947-1012.
19. Peters, J., Janzing, D., & Scholkopf, B. (2017). Elements of causal inference: Foundations and learning algorithms. MIT Press.
20. Raskutti, G., & Uhler, C. (2018). Learning directed acyclic graph models based on sparsest permutations. Stat, 7(1), e183.
21. Scutari, M., & Denis, J.-B. (2021). Bayesian networks: With examples in R (2nd ed.). CRC Press.
22. Scutari, M., Vitolo, C., & Tucker, A. (2019). Learning Bayesian networks from big data with greedy search: Computational complexity and efficient implementation. Statistics and Computing, 29, 1095-1108.
23. Steyerberg, E. W., & Harrell, F. E. (2016). Prediction models need appropriate internal, internal-external, and external validation. Journal of Clinical Epidemiology, 69, 245-247.
24. Tennant, P. W. G., Murray, E. J., Arnold, K. F., Berrie, L., Fox, M. P., Gadd, S. C., Harrison, W. J., Keeble, C., Ranker, L. R., Textor, J., Tomova, G. D., Gilthorpe, M. S., & Ellison, G. T. H. (2021). Use of directed acyclic graphs in applied health research: Review and recommendations. International Journal of Epidemiology, 50(2), 620-632.
25. Textor, J., van der Zander, B., Gilthorpe, M. S., Liskiewicz, M., & Ellison, G. T. H. (2016). Robust causal inference using directed acyclic graphs: The R package dagitty. International Journal of Epidemiology, 45(6), 1887-1894.
26. Tikka, S., & Karvanen, J. (2017). Identifying causal effects with the R package causaleffect. Journal of Statistical Software, 76(12), 1-30.
27. Van Calster, B., McLernon, D. J., van Smeden, M., Wynants, L., & Steyerberg, E. W. (2019). Calibration: The Achilles heel of predictive analytics. BMC Medicine, 17, Article 230.
28. Vowels, M. J., Camgoz, N. C., & Bowden, R. (2022). D'ya like DAGs? A survey on structure learning and causal discovery. ACM Computing Surveys, 55(4), 1-36.
29. Wynants, L., Van Calster, B., Collins, G. S., Riley, R. D., Heinze, G., Schuit, E., Bonten, M. M. J., Damen, J. A. A., Debray, T. P. A., De Vos, M., Dhiman, P., Haller, M. C., Harhay, M. O., Henckaerts, L., Lohmann, A., Jenniskens, K., Kreuzberger, N., Luijken, K., Ma, J., ... Van Smeden, M. (2020). Prediction models for diagnosis and prognosis of COVID-19: Systematic review and critical appraisal. *BMJ, 369*, Article m1328. https://doi.org/10.1136/bmj.m1328
30. Zheng, X., Aragam, B., Ravikumar, P. K., & Xing, E. P. (2018). DAGs with no tears: Continuous optimization for structure learning. Conference proceedings, 9472-9483.

XVII. APPENDICES

Appendix A. Bayesian network variable states.

Variable	States used in the Bayesian network
Sex	Female; Male
Age	Age<50; Age50-59; Age60+
Smoker	No; Yes
Cigarettes per day	Cigs0; Cigs1-10; Cigs11-20; Cigs21+
BMI	BMI<25; BMI25-29.9; BMI30+
Systolic blood pressure	SBP<120; SBP120-139; SBP140+
Total cholesterol	Chol<200; Chol200-239; Chol240+
Glucose	Glu<100; Glu100-125; Glu126+
Diabetes	No; Yes
Hypertension	No; Yes
TenYearCHD	No; Yes

Appendix B. Evidence Scenarios

Scenario	Evidence entered into the Bayesian network
Low-risk profile	Female; age <50; non-smoker; BMI <25; SBP <120; cholesterol <200; glucose <100; no diabetes; no hypertension
Older hypertension profile	Age 60+; SBP 140+; hypertension yes
Smoking-cholesterol profile	Smoker yes; cigarettes 21+; cholesterol 240+; SBP 140+
Diabetes-metabolic profile	Age 60+; BMI 30+; diabetes yes
Combined high-risk profile	Male; age 60+; smoker yes; BMI 30+; SBP 140+; cholesterol 240+; diabetes yes; hypertension yes

Appendix C. Reproducibility Checklist

Step	Reproducibility detail
Dataset audit	4,240 rows; 16 columns; no duplicate records
Missing-value handling	Mode imputation for categorical variables; median imputation for continuous variables
Discretization	Clinically interpretable categorical states
Training and testing	Stratified 70/30 split
Parameter learning	Laplace-smoothed conditional probability tables
Evaluation	AUC, average precision, Brier score, log-loss, accuracy, balanced accuracy, calibration
Interpretation	Posterior scenarios and sensitivity analysis