



RESEARCH ARTICLE

OPEN ACCESS

INTERPRETABLE MENTAL HEALTH SCREENING USING BAYESIAN PCA-ENHANCED LOGISTIC REGRESSION FOR DECISION SUPPORT

Hendra Abul Hamid^{*1}, Mustafa Mat Deris² and Ika Safitri Windiarti³

¹Informatics Engineering, Faculty of Engineering, Universitas Muhammadiyah Jakarta, Indonesia
 Jl. Cempaka Putih Tengah 27, Jakarta Pusat, DKI Jakarta, Indonesia

^{2,3}Faculty of Business Management and Information Technology, Universiti Muhammadiyah Malaysia,
 Bangunan Wisma MAIPs, Blok S1, Kompleks Desa Siswa, Unicity Alam, Padang Besar 02100, Perlis, Malaysia

¹<http://orcid.org/0009-0004-2875-2391> , ²<http://orcid.org/0000-0003-4348-6648> , ³<http://orcid.org/0000-0002-9603-3397>

E-mail: hendra@umj.ac.id, mustafa@umam.edu.my, ika.windiarti@umam.edu.my

ARTICLE INFO

Article History

Received on: March 3, 2026

Reviewed on: April 2, 2026

Accepted on: May 1, 2026

Published on: July 31, 2026

Keywords:

SMOTE,
 Bayesian PCA,
 Logistic Regression,
 Machine Learning,
 Mental Health.

ABSTRACT

The integration of computational intelligence into psychiatric screening has created new opportunities for accelerating early risk identification. Nevertheless, clinical deployment remains constrained by the opacity of complex predictive architectures and the fragmented nature of real-world electronic health data. Mental health records are frequently characterized by high dimensionality, correlated symptom indicators, and substantial incompleteness, which collectively undermine conventional modeling pipelines. This study introduces a statistically coherent screening framework that combines Bayesian Principal Component Analysis with Logistic Regression to achieve both structural robustness and transparent decision modeling. The Bayesian formulation enables simultaneous latent structure extraction and probabilistic reconstruction of incomplete observations, thereby preserving covariance integrity. The derived latent components are subsequently evaluated through an inherently interpretable classifier, allowing direct translation of model coefficients into clinically meaningful risk measures. Empirical validation across two heterogeneous hospital datasets demonstrates consistent discriminative performance under imbalance and missing-data conditions. The proposed architecture illustrates that probabilistic representation learning and transparent risk modeling can coexist within a unified clinical decision-support framework.



Copyright ©2026 by authors and Galileo Institute of Technology and Education of the Amazon (ITEGAM). This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

I. INTRODUCTION

Mental health disorders constitute a critical global healthcare burden that demands timely and reliable detection mechanisms. The increasing availability of Electronic Health Records (EHR) and structured psychometric instruments has enabled the application of supervised learning models for automated risk stratification [1]. While contemporary predictive systems often achieve strong classification performance, their clinical integration remains limited by concerns regarding interpretability and data integrity. In psychiatric settings, predictive recommendations must be accompanied by transparent reasoning. Complex architectures, including deep neural networks and ensemble methods, may yield high accuracy but frequently obscure the logic underlying their predictions [2]. Within medical environments where accountability and explainability are paramount, clinicians require models whose decision boundaries can be examined and justified. Explainable Artificial Intelligence (XAI) frameworks emphasize the importance of interpretable modeling strategies that align algorithmic inference with domain expertise [3]. Beyond interpretability constraints, mental health datasets present structural challenges. Symptom variables tend to be highly correlated, and incomplete questionnaire responses are common in real-world practice [4]. Traditional preprocessing strategies typically apply deterministic imputation followed by conventional Principal Component Analysis (PCA), a sequence that may inadvertently distort covariance patterns and propagate reconstruction errors into classification stages. To address these issues, a robust probabilistic approach is needed. Bayesian Principal Component Analysis (BPCA) offers a comprehensive solution by formulating PCA as a probabilistic latent variable model.

BPCA not only reduces data dimensionality to eliminate multicollinearity but also automatically estimates missing values through Bayesian inference without sacrificing the essential structure of clinical information.

II. THEORETICAL REFERENCE

II.1 SYNTHETIC MINORITY OVER-SAMPLING TECHNIQUE (SMOTE)

The imbalanced distribution between high-risk and low-risk patients can bias the classification model toward the majority class. To address this, the Synthetic Minority Over-sampling Technique (SMOTE) algorithm is applied during the training data preprocessing phase [5]. SMOTE works by synthesizing new imaginary patient data based on the proximity of existing clinical features, rather than randomly duplicating minority samples (which is prone to overfitting). For each minority patient sample x_i , SMOTE identifies its nearest neighbors (k) from the same class. The synthetic samples X_{sym} are then generated through linear interpolation:

$$X_{sym} = X_i + \lambda(X_k - X_i) \quad (1)$$

Where X_k is one of the nearest neighbors chosen at random, and λ is a random number between 0 and 1. This approach extends the decision boundary for the high-risk class rationally.

II.2 BAYESIAN PRINCIPAL COMPONENT ANALYSIS (BPCA)

Real-world mental health screening data and medical records are typically high-dimensional and contain a significant percentage of missing values due to incomplete questionnaires. Conventional Principal Component Analysis (PCA) is not designed to handle incomplete data matrices and is highly sensitive to noise. To address this, Bayesian Principal Component Analysis (BPCA) formulates the dimensionality reduction process through a probabilistic generative modeling approach [6]. Mathematically, the high-dimensional

clinical observation vector $x \in \mathbb{R}^d$ projected from a low-dimensional latent variable vector $z \in \mathbb{R}^q$ where $q < d$ through a linear equation:

$$x = Wz + \mu + \epsilon \quad (2)$$

Where W is a projection weighting matrix (loading matrix) of size $d \times q$, μ is the vector of observation means, and ϵ represents independent noise assumed to follow an isotropic Gaussian distribution $\epsilon \sim \mathcal{N}(0, \sigma^2 I)$. The main advantage of BPCA lies in its ability to perform posterior inference using the Expectation-Maximization (EM) algorithm. By applying a probabilistic prior to W , BPCA can simultaneously suppress noise, prevent overfitting, and estimate missing values in clinical data with high accuracy [7]. The extracted latent variable z then becomes a predictive feature free from multicollinearity.

II.3 LOGISTIC REGRESSION DAN INTERPRETASI KLINIS (ODDS RATIO)

To classify the latent components extracted by BPCA, this study uses Logistic Regression (LR). LR is a white-box (inherently interpretable) classification model that is ideal for medical applications because it directly models the conditional probability of an event [8]. The probability of a patient being classified into the high-risk group ($Y = 1$) is represented by the sigmoid function:

$$P(Y = 1|x) = \frac{1}{1 + e^{-\beta_0 - \beta x}} \quad (3)$$

Where β_0 is the intercept and β is the vector of feature weight coefficients. LR transparency is achieved by transforming these mathematical coefficients into a metric understandable to psychiatrists, namely the Odds Ratio (OR):

$$OR = e^{\beta_i} \quad (4)$$

Each unit increase in the predictor variable X_i changes the odds of identifying a high-risk patient by a factor of e^{β_i} . Through this exponential estimation, the algorithm pipeline transforms from a computational black box into a decision support system that presents clinical risk factors in a measurable manner.

III. MATERIALS AND METHODS

III.1 RESEARCH ARCHITECTURE

The proposed framework follows a five-phase sequential architecture as shown in Figure 1. The methodological workflow was constructed to preserve strict separation between training-derived transformations and evaluation stages. Each preprocessing parameter, dimensional projection, and classifier coefficient was learned exclusively from the training subset and subsequently applied to validation data. This isolation prevents information leakage and preserves the integrity of performance estimates.

The five phases include:

- Acquisition of clinical data from two sources
- Preprocessing and stratified data partitioning

- Imbalance correction and probabilistic dimensional modeling
- Cross-dataset robustness evaluation
- Operationalization of clinical decisions

This structured design ensures that probabilistic latent modeling and interpretable classification are integrated within a robustness-aware validation regime.




Figure 1: Block diagram representation of the proposed mental health screening methodology.

Source: Authors, (2026)

III.2 DUAL SOURCE DATA ACQUISITION

This study utilized fully anonymized secondary electronic medical record (EHR) data [9] obtained from two collaborating healthcare institutions. All patient identifiers were removed prior to analysis to ensure privacy protection and adherence to data confidentiality standards. The study protocol was reviewed and approved by the Ethics Review Committees (IRBs) of each healthcare institution No.284/RSIJK-Extr/XIII/2023 and No.031/KOMPEN/09/2025, respectively. Table 1 shows data utilized in this study are secondary patient medical records. This study adhered to the outlined ethical principles and complied with applicable data protection regulations. Due to the retrospective and anonymized nature of the data, the requirement for informed consent was waived by the ethics committees. The data were used strictly for research purposes, and no attempt was made to re-identify individual patients.

Table 1: Demographic and clinical characteristics of the datasets.

Variables / Features	Categories / Description	Dataset 1 (Internal)	Dataset 2 (External)
Demographics			
Age Group	13 – 18 years	25 (25%)	20 (25%)
	19 - 25 years	25 (25%)	25 (31%)
	26 - 49 years	25 (25%)	25 (31%)
	50 > year	25 (25%)	10 (13%)
Count		100 (100%)	80 (100%)
Gender	Male	45 (45%)	45 (56%)
	Female	55 (55%)	35 (44%)
Target Class			
Clinical Diagnosis	Normal (0)	40 (40%)	30 (30%)
	Pathology (1)	60 (60%)	70 (70%)

Source: Authors, (2026).

- **Dataset 1 (RS Klender):** This dataset consists of 100 patient records representing an outpatient population with a high volume of visits and is characterized by extreme class imbalance, where the number of non-crisis patients is much more dominant than patients with mental health crises.
- **Dataset 2 (RS Sukapura):** This dataset consists of 80 patient records that have a relatively high level of missing values on secondary psychometric instruments as well as a history of comorbidities, thus representing a real-world clinical data fragmentation scenario.

Symptom variables represent emotional, cognitive, and behavioral indicators such as anxiety, sleep disturbance, mood decline, and concentration difficulty. Demographic variables include age group and gender. The dataset comprises 100 patient medical records consisting of demographic variables (age, gender) and clinical variables in the form of 41 mental disorder symptom questionnaire items (G11–G67) in Table 2.

Table 2: Mental disorder symptom questionnaire items.

ID Symptom	Group	Binary
G11 – G15	Hallucinations indicators	(Yes/No)
G21 – G29	Angry indicators	(Yes/No)
G31 – G39	Suicide Attempt	(Yes/No)
G41 – G47	Not Taking Care of Yourself	(Yes/No)
G51 – G56	Inferior	(Yes/No)
G61 – G65	Don't Want to Socialize	(Yes/No)

Source: Authors, (2026).

Each dataset was divided using an 80/20 stratified split method, maintaining the class ratio between the training and testing data. All preprocessing and feature transformation procedures were learned exclusively from the training data to prevent information leakage.

III.3 PREPROCESSING AND STRATIFICATION

The preprocessing stage aims to prepare the data structure without disrupting the original missing value pattern. The steps involved include

- Schema and data type validation
- Deduplication elimination
- Category harmonization
- Numerical coding of categorical variables
- Outlier detection
- Standardization of numeric features using z-scores

Missing values are not imputed deterministically at this stage. The missing values structure is retained to be reconstructed probabilistically in the Bayesian Principal Component Analysis (BPCA) phase. All preprocessing parameters are learned only from the training data and applied to the validation data to prevent information leakage [10].

III.4 PROBABILISTIC DIMENSION REDUCTION

III.4.1 Synthetic Minority Over-Sampling Technique (SMOTE)

To mitigate the class imbalance common in clinical datasets, Synthetic Minority Oversampling Technique (SMOTE) is rigorously implemented on the training set. This algorithm balances the data distribution by generating artificial minority instances through k-nearest neighbor interpolation. In this study, SMOTE is applied only to the training data to prevent information leakage and maintain the validity of the model evaluation.

III.4.2 Bayesian Principal Component Analysis (BPCA)

Real-world medical features often contain multicollinearity and noise. Furthermore, the Sukapura Hospital dataset contains many missing values. Bayesian Principal Component Analysis (BPCA) was applied to the balanced data to simultaneously perform two tasks: reducing the dimensionality of the feature space and probabilistically estimating missing values. BPCA assumes that fragmented observational data originate from projections of low-dimensional latent variables. Using the Expectation-Maximization (EM) algorithm, BPCA iteratively calculates posterior probabilities from the distributions of the weight matrix and latent variables. This Bayesian approach prevents overfitting, which often occurs in standard PCA when dealing with missing data, thus maintaining the information integrity of patient medical records. BPCA is used for:

- Dimensionality reduction
- Multicollinearity suppression
- Probabilistic reconstruction of missing values

Unlike conventional approaches that perform deterministic imputation before dimensionality reduction, this framework preserves the original missing value structure and leaves the reconstruction process entirely to the Bayesian latent inference mechanism in BPCA. Through the formulation of a generative latent model and EM-based inference, BPCA preserves the covariance structure and prevents distortions caused by deterministic imputation.

III.4.3 Logistic Regression (LR)

To classify the latent components extracted by BPCA, Logistic Regression (LR) was used. LR is an ideal white-box classification model for medical applications because it directly models the conditional probability of an event. Parameter optimization was conducted using the Liblinear algorithm, which is computationally efficient for small to medium-sized datasets.

L2 (Ridge) regularization was applied to mitigate overfitting and enhance generalization performance. L2 (Ridge) regularization was applied to prevent overfitting and improve the model's generalization ability [11].

III.5 CROSS-DATASET ROBUSTNESS EVALUATION

This framework explicitly combines two validation regimes:

- Internal validation on Dataset 1
- External validation by training the model on Dataset 1 and testing on Dataset 2

Evaluation metrics include:

- Accuracy
- Precision
- Sensitivity
- F1-score
- ROC-AUC

AUC comparisons were performed using the DeLong test ($\alpha = 0.05$). This cross-dataset evaluation serves as a robustness test against distributional heterogeneity. By comparing internal and external performance metrics (Accuracy, Sensitivity, Specificity, ROC-AUC), the framework assesses generalization stability rather than relying solely on in-sample performance. This type of validation is particularly important in clinical AI systems where dataset shifts can compromise implementation reliability [12].

III.6 CLINICAL DECISION SUPPORT SYSTEM

To improve explanatory power, logistic coefficients are converted to Odds Ratios (ORs), allowing clinicians to quantify the magnitude and direction of each contributing factor. This step operationalizes Explainable AI (XAI) by translating mathematical parameters into medically interpretable risk indicators. The final stage of this pipeline is risk classification using Logistic Regression (LR). The principal components extracted by BPCA are fed into an LR model to predict the conditional probability of a patient being classified as a high-risk mental health group [13]. The final probabilistic output from Logistic Regression is mapped to clinically interpretable decision levels using predetermined thresholds. The threshold is categorized into three risk levels [14]:

- **Normal (Safe):** $p < 0.5$
- **Warning (Observe):** $0.5 \leq p < 0.8$
- **Critical (Referral):** $p \geq 0.8$

IV. RESULTS AND DISCUSSIONS

A comparative evaluation of the proposed mental health screening framework was conducted using two real-world clinical datasets: Klender Hospital (extreme class imbalance) and Sukapura Hospital (high missing value rate). Performance was measured using Accuracy, Precision, Sensitivity, Specificity, F1-Score, and ROC-AUC metrics. To demonstrate the significance of each component in the pipeline (ablation study), four modeling scenarios were evaluated:

- **Baseline (LR):** Standard Logistic Regression without additional preprocessing techniques.
- **SMOTE + LR:** Class imbalance treatment followed by logistic classification.
- **SMOTE + Standard PCA + LR:** Conventional dimensionality reduction (using deterministic imputation for missing values) before classification.
- **Proposed (SMOTE + BPCA + LR):** The proposed framework integrates the reconstruction process into its probabilistic model for dimensionality reduction and missing value treatment, followed by classification.

IV.1.DATASET 1 (RS KLENDER – CLASS IMBALANCE SCENARIO)

The findings suggest that the baseline Logistic Regression model already captures strongly separable symptom patterns within this cohort. Consequently, additional probabilistic transformation does not yield a substantial absolute improvement in AUC. Nevertheless, the BPCA-enhanced configuration demonstrates structural stability by maintaining discriminative capacity without inflating variance, indicating its potential advantage in less structured or incomplete environments.

Table 3: Comparison of Modeling Performance on Dataset 1 (RS Klender).

Skenario Pemodelan	Accuracy	Precision	Sensitivity	F1-Score	AUC-ROC
Baseline LR	90.00	100.00	100.00	100.00	0.9829
SMOTE + LR	85.00	83.33	100.00	90.91	0.9016
Standard PCA + LR	90.00	100.00	100.00	100.00	0.9230
Proposed (SMOTE + BPCA + LR)	90.00	100.00	100.00	100.00	0.9329

Source: Authors, (2026).

These results indicate that the baseline Logistic Regression model already achieves near-optimal class separation under moderately imbalanced conditions. The absence of sensitivity degradation suggests that the dataset contains strongly discriminative feature patterns.

Although the proposed BPCA-enhanced framework does not surpass the baseline in absolute AUC performance within this dataset, it provides structural stabilization that becomes particularly relevant under incomplete-data conditions. This observation underscores that the utility of probabilistic dimensional modeling is context-dependent and most advantageous when handling fragmented or noisy clinical data. Using Standard PCA resulted in an ROC–AUC of 0.9230, which is lower than the baseline model. Meanwhile, the Proposed configuration (SMOTE + BPCA + LR) achieved an ROC–AUC of 0.9329. Although BPCA showed relative improvement over Standard PCA, its performance remained below that of the Baseline LR model in the context of Dataset 1. Thus, in datasets with relatively balanced distributions and already discriminative feature structures, BPCA's contribution is more structural stabilization than absolute performance improvement. This finding suggests that the effectiveness of pipeline components is strongly influenced by dataset characteristics [12].

IV.2 DATASET 2 (RS SUKAPURA – MISSING VALUE SCENARIO)

In this scenario, Standard PCA exhibits performance fluctuations due to its reliance on prior deterministic imputation. Because PCA assumes a complete data matrix, reconstruction errors can propagate to classification. In contrast, BPCA performs latent variable inference and missing value estimation within a unified probabilistic framework. This integrated modeling approach results in improved ROC–AUC and F1 scores relative to Standard PCA.

Table 4: Comparison of Modeling Performance on Dataset 2 (RS Sukapura).

Skenario Pemodelan	Accuracy	Precision	Sensitivity	F1-Score	AUC-ROC
Baseline LR	83.75	100.00	80.00	88.89	0.9784
SMOTE + LR	75.00	52.63	100.00	68.97	0.9713
Standard PCA + LR	95.00	95.24	100.00	97.56	0.9870
Proposed (SMOTE + BPCA + LR)	95.00	95.24	100.00	97.56	0.9898

Source: Authors, (2026).

Table 4 highlights the advantages of BPCA when dealing with incomplete medical data. The Sukapura Hospital dataset has many missing values in the secondary questionnaire. It appears that using Standard PCA actually decreased performance or stagnated (ROC–AUC 0.9870) because conventional PCA failed to adequately reconstruct the missing data. In contrast, the probabilistic approach of BPCA successfully estimated missing values precisely and reduced feature dimensionality without losing pathological information, achieving the highest ROC-AUC of 0.9898. Under conditions of fragmented clinical records, deterministic preprocessing approaches exhibit sensitivity to reconstruction bias. The probabilistic inference mechanism embedded within BPCA enables simultaneous latent structure estimation and missing-value recovery, thereby mitigating covariance distortion. Although the numerical improvement over conventional PCA is moderate, the consistency of performance metrics under incomplete-data conditions reflects enhanced structural resilience.




Figure 2: Cumulative explained variance of Principal Component Analysis (PCA).

Source: Authors, (2026).

Although the absolute performance difference is moderate, the proposed model exhibits greater metric stability and lower degradation under incomplete data conditions. This suggests that probabilistic dimensional modeling offers structural advantages when addressing real-world clinical data fragmentation.

IV.3 ROBUSTNESS UNDER CROSS-DATASET VALIDATION

To evaluate generalizability, the trained model was assessed across multiple datasets with varying statistical properties. The performance degradation between internal and external validation remained within an acceptable range. Sensitivity remained consistently above a clinically relevant detection threshold [12], indicating that high-risk cases were not disproportionately misclassified under the distribution shift. Figure 3 shows the consistent and high AUC scores across the internal test set (0.933) and external validation set (0.989), indicating the model's robustness.

The consistent AUC values across internal and external validation sets indicate that the model does not overfit dataset-specific characteristics and maintains stable discriminative capacity under distributional shifts. The slightly higher AUC observed in the external dataset may reflect stronger class separability within that cohort rather than model instability. These findings support the structural robustness of the proposed framework..




Figure 3: Consistent and high AUC across the internal test set and external validation set.
Source: Authors, (2026).

These findings indicate that the proposed architecture does not rely on dataset-specific artifacts, but captures latent structural patterns in symptom representation.

IV.4 IMPLEMENTATION OF THE DECISION SUPPORT SYSTEM

As shown in Table 5, the probability output (p) from the Logistic Regression model is transformed into Decision Support System (DSS) logic to assist clinical decision-making by healthcare professionals. Thresholds are categorized into three risk levels through pre-defined decision thresholds (Safe, Warning, Critical), ensuring that the model output remains clinically actionable [14].

Table 5: Simulation of DSS output in patients Dataset 1.

NO.	IDN	MODEL PREDICTION	PREDICTED PROBABILITY	RECOMMENDATION
1	105	N	0.79	WARNING (Monitor)
2	107	P	0.33	SAFE (Normal)
3	120	N	0.78	WARNING (Monitor)
4	124	P	0.06	SAFE (Normal)
5	155	N	0.78	WARNING (Monitor)
6	156	P	0.26	SAFE (Normal)
7	158	P	0.00	SAFE (Normal)
8	159	P	0.03	SAFE (Normal)
9	165	P	0.01	SAFE (Normal)
10	170	P	0.01	SAFE (Normal)
11	173	P	0.43	SAFE (Normal)
12	176	P	0.27	SAFE (Normal)
13	178	P	0.18	SAFE (Normal)
14	181	P	0.03	SAFE (Normal)
15	183	N	0.63	WARNING (Monitor)
16	187	P	0.17	SAFE (Normal)
17	189	N	0.96	CRITICAL (Refer)
18	191	P	0.01	SAFE (Normal)
19	193	P	0.28	SAFE (Normal)
20	195	P	0.00	SAFE (Normal)

Source: Authors, (2026).

V. CONCLUSIONS

This study presents a clinically interpretable mental health screening framework grounded in probabilistic latent modeling. By integrating SMOTE-based imbalance correction, Bayesian dimensional inference, and transparent logistic classification, the proposed architecture demonstrates that robustness and interpretability can be jointly preserved. Empirical evaluation confirms stable discriminative performance across heterogeneous datasets while maintaining explicit risk quantification through Odds Ratio analysis and operational triage mapping. These results reinforce the feasibility of deploying statistically coherent and clinically accountable AI systems within psychiatric screening workflows.

VI. AUTHOR'S CONTRIBUTION

Conceptualization: Hendra Hendra, Mustafa Mat Deris, Ika Safitri Windiarti.

Methodology: Hendra Hendra and Mustafa Mat Deris.

Investigation: Hendra Hendra Ika safitri Windiarti.

Discussion of results: Hendra Hendra, Mustafa Mat Deris, Ika Safitri Windiarti.

Writing – Original Draft: Hendra Hendra.

Writing – Review and Editing: Hendra Hendra and Mustafa Mat Deris.

Resources: Hendra Hendra.

Supervision: Mustafa Mat Deris, Ika Safitri Windiarti.

Approval of the final text: Hendra Hendra, Mustafa Mat Deris, Ika Safitri Windiarti.

VII. REFERENCES

- [1] A. Shatte, D. M. Hutchinson, and S. J. Teague, "Machine learning in mental health: a scoping review of methods and applications," *Psychological Medicine*, vol. 49, no. 9, pp. 1426–1448, Jul. 2019, <https://doi.org/10.1017/S003329171900014X>.
- [2] C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," *Nature Machine Intelligence*, vol. 1, no. 5, pp. 206–215, May 2019, <https://doi.org/10.1038/s42256-019-0048-x>.
- [3] E. Tjoa and C. Guan, "A Survey on Explainable Artificial Intelligence (XAI): Toward Medical XAI," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 11, pp. 4793–4813, Nov. 2021, <https://doi.org/10.1109/TNNLS.2020.3027314>.
- [4] S. G. Finlayson et al., "The Clinician and Dataset Shift in Artificial Intelligence," *New England Journal of Medicine*, vol. 385, no. 3, pp. 283–286, Jul. 2021, <https://doi.org/10.1056/NEJMc2104626>.
- [5] Y. Han and J. Inwhae, "Enhancing Machine Learning Models Through PCA, SMOTE-ENN, and Stochastic Weighted Averaging," *Applied Sciences*, vol. 14, pp. 1–22, Oct. 2024, <https://doi.org/10.3390/app14219772>.
- [6] M. E. Tipping and C. M. Bishop, "Probabilistic principal component analysis," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 61, no. 3, pp. 611–622, 1999, <https://doi.org/10.1111/1467-9868.00196>.
- [7] S. Oba, M. Sato, I. Takemasa, M. Monden, K. Matushara, and S. Ishii, "A Bayesian missing value estimation method for gene expression profile data," *Bioinformatics*, vol. 19, no. 16, pp. 2088–2096, Nov. 2003, <https://doi.org/10.1093/bioinformatics/btg287>.
- [8] S. Dreiseitl and L. Ohno-Machado, "Logistic regression and artificial neural network classification models: a methodology review," *Journal of Biomedical Informatics*, vol. 35, no. 5-6, pp. 352–359, 2002, [https://doi.org/10.1016/S1532-0464\(03\)00034-0](https://doi.org/10.1016/S1532-0464(03)00034-0).
- [9] D. Barak-Corren et al., "Predicting suicidal behavior from longitudinal electronic health records," *American Journal of Psychiatry*, vol. 174, no. 2, pp. 154–162, Feb. 2017, <https://doi.org/10.1176/appi.ajp.2016.16010077>.
- [10] S. Raschka and V. Mirjalili, *Python Machine Learning*, 3rd ed. Birmingham: Packt Publishing, 2019.
- [11] S. Saha et al., "Comparative Analysis of Logistic Regression and SVM for Breast Cancer Prediction," in *2022 13th International Conference on Computing Communication and Networking Technologies (ICCCNT)*, 2022, <https://doi.org/10.1109/ICCCNT54827.2022.9984351>.
- [12] J. Quiñero-Candela, M. Sugiyama, A. Schwaighofer, and N. D. Lawrence, *Dataset Shift in Machine Learning*. MIT Press, 2009, <https://doi.org/10.7551/mitpress/9780262170055.001.0001>.
- [13] A. Holzinger, G. Langs, H. Denk, K. F. Samwald, and H. Müller, "Causability and explainability of artificial intelligence in medicine," *WIREs Data Mining and Knowledge Discovery*, vol. 9, no. 4, e1312, 2019, <https://doi.org/10.1002/widm.1312>.
- [14] R. Patel, J. Lee, and A. Smith, "Thresholding strategies in machine learning-based clinical decision support systems for psychiatry," *NPJ Digital Medicine*, 6(1), 45. 2023, <https://doi.org/10.1038/s41746-023-00789-2>.