

A Lightweight Hybrid Machine Learning Framework for Parameterized Human Behavior Classification

Dr Jyoti Bala Gupta¹, Brijesh Kumar², Ashish Kumar Pandey³

^{1,2,3}Department of Computer Science & Information Technology, Dr. C. V. Raman University, Bilaspur, Chhattisgarh, India

Abstract:

Human-behavior modeling is commonly approached through video, sensor, handwriting, or textual signals, yet many practical settings require a lightweight model that can operate on structured behavioral responses. This study develops a hybrid supervised-learning framework that combines logistic regression with transparent rule-based decision logic for binary classification of parameterized behavioral-response patterns. The experimental dataset described in the source study contains 500 questionnaire-derived instances represented by ten behavioral attributes, including responsibility, gratitude, patience and listening, positive attitude, respectfulness, confidence, humility, kindness and compassion, honesty, and justice and integrity. A stratified 90:10 train-test protocol was used, supplemented by five-fold cross-validation. The source confusion matrix for the 50-instance hold-out set contains 24 true negatives, 24 true positives, one false positive, and one false negative, corresponding to 96% hold-out accuracy, precision, recall, F1-score, and specificity. Five-fold cross-validation is reported at $97.00 \pm 0.63\%$, while ROC-AUC is reported as 0.985. The proposed approach is computationally simple and interpretable, but its outputs must be treated as probabilistic decision-support signals rather than deterministic judgments about individuals. The study therefore emphasizes human oversight, privacy safeguards, external validation, and bias auditing before any real-world deployment.

Keywords: human behavior analysis, logistic regression, rule-based classification, supervised learning, interpretable machine learning, behavioral questionnaire

I. INTRODUCTION

Automated analysis of human activity and behavior has evolved from manual observation toward computational approaches that use vision, wearable sensors, textual information, and structured signals. Earlier work has shown that machine learning can support human activity recognition and behavioral analysis, but many high-performing approaches depend on video streams, large neural networks, or sensor infrastructure. Such requirements can reduce transparency and make deployment difficult in resource-constrained settings [1]–[5].

The present study investigates a compact alternative: a structured questionnaire converted into ten numerical behavioral features and evaluated using a hybrid of Logistic Regression and an explicit rule-based layer. The source study frames the task as binary supervised classification. For publication-quality reporting, the two classes are described here as prosocial response patterns and anomalous/high-risk

response patterns rather than as absolute labels of a person's character. This distinction is important because questionnaire-based machine-learning outputs are statistical estimates and cannot establish intent, morality, criminality, or psychological diagnosis.

The main contributions are: (i) a lightweight ten-feature behavioral representation; (ii) an interpretable hybrid classification pipeline; (iii) hold-out and cross-validation reporting; (iv) comparison with common baseline classifiers; and (v) explicit discussion of privacy, bias, and human oversight requirements.

II. RELATED WORK

Human behavior analysis spans several modalities. Pantic et al. surveyed human computing and machine understanding of behavior and highlighted the need to connect observable cues with computational interpretation [1]. Popoola and Wang reviewed abnormal behavior recognition from video and summarized challenges in scene understanding and event recognition [2]. More recent surveys have examined sensor-based and vision-based human activity recognition, noting trade-offs among feature engineering, deep architectures, computational cost, and real-time operation [3].

Deep learning has improved representation learning for activity recognition. CNN-LSTM combinations are especially useful when spatial information and temporal dynamics must be modeled jointly [4], [5]. Neuro-fuzzy and expert-informed models have also been proposed to combine data-driven learning with human knowledge [6]. In organizational and predictive settings, machine learning has been discussed as a way to extract behavior-related patterns from diverse data streams [7], [8].

A recurring gap is that many systems focus on physical actions, facial expressions, or sensor trajectories rather than a compact and transparent questionnaire-based representation. The present work therefore studies whether a small set of structured behavioral attributes can support a reproducible binary classification baseline without requiring specialized sensing hardware.

A. Positioning of the Present Study

The source literature spans four broad strategies: visual action recognition, wearable or ambient sensing, text/opinion analysis, and expert-knowledge systems. Vision and sensor systems are valuable when movement is the target, but they require cameras, devices, or continuous signal acquisition. Text-based methods capture expressed opinions, while expert-informed fuzzy systems can encode domain rules. In contrast, the present study intentionally uses a small questionnaire-derived feature vector. Its goal is not to replace multimodal behavioral science, but to establish a low-cost and explainable baseline that can be reproduced with standard machine-learning software.

This positioning also clarifies the meaning of the target variable. The model classifies response patterns defined by the dataset construction process; it does not directly observe a person's internal mental state. Consequently, performance metrics quantify agreement with the supplied labels rather than objective truth about personality. This distinction is central to responsible interpretation of the results.

III. MATERIALS AND METHODS

A. Dataset and Behavioral Features

The source study reports a primary dataset of 500 respondents aged 18 years and above, collected through a structured questionnaire and observational experience. Each record is represented through ten input attributes. The attributes are encoded numerically and used as the feature vector X . The target is binary. Because the original terminology uses value-laden labels, this paper uses the operational labels

Class 0 (prosocial response pattern) and Class 1 (anomalous/high-risk response pattern) solely for model evaluation.

The ten attributes used in the source study are: (1) Responsibility, (2) Gratitude, (3) Patience and Listening, (4) Positive Attitude, (5) Respectfulness, (6) Confidence, (7) Humility, (8) Kindness and Compassion, (9) Honesty, and (10) Justice and Integrity. The source document describes these as binary or numerically encoded questionnaire responses.

B. Preprocessing and Data Split

The reported preprocessing pipeline includes missing-value handling, numerical encoding, feature scaling, and stratified partitioning. The principal experiment uses a 90:10 train-test split, giving 450 training samples and 50 hold-out samples. The source also reports five-fold cross-validation over the complete dataset to assess stability across partitions.

C. Hybrid Classification Model

The predictive model combines a transparent rule-based screening stage with Logistic Regression. The rule layer handles explicit response-pattern conditions, while Logistic Regression estimates a probabilistic decision boundary from the ten input variables. For an input vector x , the logistic model estimates the probability of Class 1 as

$$p(y=1|x) = 1 / (1 + \exp(-(w^T x + b)))$$

where w is the learned weight vector and b is the intercept. A decision threshold converts the posterior probability into the final class. The rule-based component is retained as an interpretable decision-support layer rather than a replacement for statistical validation.




Fig. 1. Proposed lightweight hybrid classification workflow.

D. Evaluation Measures

Performance was assessed using accuracy, precision, recall (sensitivity), F1-score, specificity, ROC-AUC, confusion-matrix analysis, and five-fold cross-validation. These measures were selected to avoid relying on accuracy alone and to expose false-positive and false-negative behavior.

E. Reproducibility Protocol

The implementation environment described in the source study uses Python 3.10 with Pandas, NumPy, Scikit-Learn, and Matplotlib. A reproducible execution should preserve the same feature order, preprocessing rules, random seed, stratified split, threshold, and cross-validation folds. The source also lists an ordinary desktop-class configuration (Intel Core i5/AMD Ryzen 5 class processor, 8–16 GB RAM), indicating that the method does not depend on GPU acceleration. For a final archival version, the authors should preserve the de-identified dataset, the preprocessing script, and the exact model parameters used to generate the reported tables and figures.

IV. RESULTS

A. Hold-out Performance

The source document reports a 50-sample hold-out test set and provides a confusion matrix with $TN = 24$, $FP = 1$, $FN = 1$, and $TP = 24$. These counts imply 48 correct classifications out of 50, i.e., 96% hold-out accuracy. Precision, recall, F1-score, and specificity are also 96% for the balanced 25/25 test subset.




Fig. 2. Confusion matrix reported for the 50-sample hold-out set.

Hold-out metrics derived from the reported confusion matrix are: Accuracy = 96.00%, Precision = 96.00%, Recall/Sensitivity = 96.00%, F1-score = 96.00%, and Specificity = 96.00%. The source study additionally reports ROC-AUC = 0.985 and five-fold cross-validation accuracy = $97.00 \pm 0.63\%$.

B. Cross-Validation

Five-fold cross-validation is reported with fold accuracies of 97%, 96%, 98%, 97%, and 97%, giving a mean of 97.00% and a reported dispersion of $\pm 0.63\%$. This result is more informative than a single train-test split because it evaluates model stability across multiple partitions. Nevertheless, cross-validation on the same source dataset does not substitute for external validation on independently collected data.

C. Error Analysis

The confusion matrix contains one false positive and one false negative. In a behavioral decision-support setting, the two error types have different implications. A false positive may incorrectly flag a prosocial response pattern, while a false negative may fail to flag an anomalous pattern. Because each error can affect individuals differently, operational use should not rely on a single threshold chosen only to maximize accuracy. Threshold selection should consider calibration, context, and the cost of each error type, and any consequential action should require independent human verification.

D. Comparison with Baseline Models

The source study also reports comparisons with Naive Bayes, Linear SVM, and Decision Tree classifiers. The proposed hybrid approach has the strongest reported cross-validated performance, while the confusion-matrix-consistent hold-out accuracy is 96%. The comparison is presented as an exploratory benchmark because the dataset is relatively small and the source document does not provide complete hyperparameter specifications for every baseline.




Fig. 3. Hold-out accuracy comparison using source-reported baseline values and confusion-matrix-derived hybrid accuracy.

V. DISCUSSION

The principal advantage of the proposed framework is its simplicity. Ten structured features require little computational overhead, Logistic Regression is directly interpretable through learned coefficients, and a rule layer can express domain constraints in a human-readable form. Such characteristics may be useful when the objective is to build a transparent baseline before moving to more complex multimodal models.

At the same time, the strong numerical results should be interpreted cautiously. A high score on a curated questionnaire dataset may reflect regularities in the labeling process as much as generalizable behavioral structure. The model does not observe context, temporal change, culture, language, situational stress, or disagreement among human raters. Therefore, it should not be used to declare that an individual is intrinsically 'good' or 'bad,' to infer criminal intent, or to make consequential decisions without independent evidence and human review.

The reported training accuracy of 100% also warrants additional validation. Although the cross-validation results are stable, a larger external dataset is needed to assess overfitting, demographic robustness, calibration, and fairness. Future evaluations should report class distribution, inter-rater agreement for target labels, questionnaire reliability, feature definitions, and confidence intervals.

VI. PRACTICAL IMPLICATIONS

The framework is most defensible as a research prototype for studying whether structured self-report or assessor-coded attributes contain learnable patterns. In education, organizational studies, or counseling research, it could be used to explore aggregate trends when participation is voluntary and data are anonymized. It may also serve as a teaching example of how interpretable machine learning differs from opaque end-to-end models.

The framework is not validated for autonomous screening in employment, law enforcement, admissions, discipline, or other high-stakes settings. Such uses would require independent prospective evaluation, documented construct validity, fairness testing, legal review, and a clear mechanism for people to challenge or correct automated inferences. The present dataset and reported evaluation are insufficient for those purposes.

VII. ETHICAL, PRIVACY AND FAIRNESS CONSIDERATIONS

Behavioral questionnaires involve human-participant data and can produce sensitive inferences. Any deployment should therefore include informed consent, data minimization, de-identification, secure

storage, restricted access, and a documented retention policy. The journal's author guidance also requires appropriate institutional approval for human-participant research and a statement that informed consent was obtained. These items must be confirmed by the authors before submission.

Model outputs should be treated as advisory probabilities. A human-in-the-loop process should allow review and appeal, and periodic bias audits should test whether error rates differ across demographic or socio-economic groups. High-stakes use in employment, policing, education, credit, or legal settings requires much stronger validation than is provided by the present exploratory dataset.

VIII. LIMITATIONS AND FUTURE WORK

- The dataset contains 500 instances and requires validation on larger, independently collected samples.
- The ten features are simplified representations of behavior and do not capture duration, context, intensity, or longitudinal change.
- The labeling methodology, annotator agreement, and questionnaire reliability should be documented in greater detail.
- External validation, probability calibration, fairness analysis, and ablation studies are needed before operational use.
- Future extensions may compare ensemble methods and neural models, and may integrate multimodal inputs such as audio or vision only when consent and privacy safeguards are in place.

IX. CONCLUSION

This paper presents a compact hybrid machine-learning baseline for parameterized human-behavior response classification. Using ten questionnaire-derived features and 500 reported samples, the framework combines Logistic Regression with rule-based logic and is evaluated using a 90:10 hold-out split and five-fold cross-validation. The source confusion matrix yields 96% hold-out accuracy with balanced precision and recall, while the reported cross-validation mean is $97.00 \pm 0.63\%$ and ROC-AUC is 0.985. The model demonstrates that structured behavioral responses can be classified efficiently with an interpretable pipeline; however, the findings should be viewed as exploratory. Reliable real-world use requires ethical approval, consent, external datasets, detailed labeling protocols, fairness testing, and human oversight.

DECLARATIONS

Ethics approval and consent to participate: [MANDATORY BEFORE SUBMISSION: insert the name of the approving institutional ethics committee/authority and approval number, or provide a justified exemption statement. Confirm that informed consent was obtained from participants.]

Data availability: The source document describes a primary questionnaire dataset of 500 records. No public repository identifier is provided in the supplied material. Authors should state how a de-identified dataset can be accessed, where ethically and legally permissible.

Funding: [Insert funding statement or "This research received no external funding," if accurate.]

Conflict of interest: [Authors to confirm and insert final declaration.]

REFERENCES

1. M. Pantic, A. Pentland, A. Nijholt, and T. S. Huang, "Human computing and machine understanding of human behavior: A survey," *Lecture Notes in Computer Science*, vol. 4451, pp. 47–71, 2007, doi: 10.1007/978-3-540-72348-6_3.
2. O. P. Popoola and K. Wang, "Video-based abnormal human behavior recognition—a review," *IEEE Trans. Syst., Man, Cybern. C*, vol. 42, no. 6, pp. 865–878, 2012, doi: 10.1109/TSMCC.2011.2178594.
3. L. Minh Dang, K. Min, H. Wang, M. J. Piran, C. H. Lee, and H. Moon, "Sensor-based and vision-based human activity recognition: A comprehensive survey," *Pattern Recognition*, vol. 108, 107561, 2020, doi: 10.1016/j.patcog.2020.107561.
4. I. U. Khan, S. Afzal, and J. W. Lee, "Human activity recognition via hybrid deep learning based model," *Sensors*, vol. 22, no. 1, 323, 2022, doi: 10.3390/s22010323.
5. P. Sabitha, S. Sai, and S. Abhishek, "Human activity recognition using CNN & LSTM," in *Proc. 5th Int. Conf. Information Technology Research (ICITR)*, 2020, pp. 386–390, doi: 10.1109/ICITR51448.2020.9310792.
6. G. Mohamed, A. Lotfi, and A. Pourabdollah, "Human activities recognition based on neuro-fuzzy finite state machine," *Technologies*, vol. 6, no. 4, 110, 2018, doi: 10.3390/technologies6040110.
7. E. Real de Oliveira and P. Rodrigues, "A review of literature on human behaviour and artificial intelligence: Contributions towards knowledge management," *Electronic Journal of Knowledge Management*, vol. 19, no. 2, pp. 165–179, 2021, doi: 10.34190/ejkm.19.2.2459.
8. T. Sajana and M. Gulhane, "Human behavior prediction and analysis using machine learning—A review," *Turkish Journal of Computer and Mathematics Education*, vol. 12, no. 5, pp. 870–876, 2021, doi: 10.17762/turcomat.v12i5.1499.
9. H. Mu, R. Sun, G. Yuan, and Y. Wang, "Abnormal human behavior detection in videos: A review," *Information Technology and Control*, vol. 50, no. 3, pp. 522–545, 2021, doi: 10.5755/j01.itc.50.3.27864.
10. Y. Kim, M. Imani, and T. S. Rosing, "Efficient human activity recognition using hyperdimensional computing," in *Proc. ACM/IEEE Design Automation Conf. workshops*, 2018, doi: 10.1145/3277593.3277617.
11. S. K. S. and A. Danti, "Human behavior prediction based on opinions using machine learning techniques," *International Journal of Recent Technology and Engineering*, vol. 8, no. 6, pp. 3117–3120, 2020, doi: 10.35940/ijrte.F8733.038620.
12. B. Nguyen, Y. Coelho, T. Bastos, and S. Krishnan, "Trends in human activity recognition with focus on machine learning and power requirements," *Machine Learning with Applications*, vol. 5, 100072, 2021, doi: 10.1016/j.mlwa.2021.100072.