Natural Resources for Human Health
eISSN: 2583-1194

Original Research

Received 04 May 2026
Revised 25 June 2026
Accepted 10 July 2026

Natural Resources for Human Health 2026, Vol. 6 (12s): 708–718

KEYWORDS:
Yoga; Machine learning; Cognitive function; Evidence mapping; Systematic review protocol; Predictive modelling

Impact of Yoga Practices on Cognitive Functions through a Machine Learning Approach: A Protocol for Evidence Mapping and Predictive Modelling

Jeevan K P[1]*, Shobha S[2], P Sandhya[3]

[1,2,3] Dept. of CSE, Visvesvaraya Technological University, Belagavi, VTU CPGS Mysuru, Karnataka, India
mail: jeevankphsn@gmail.com[1] · sshobha233@gmail.com [2,] Sanjoshi17@yahoo.com [3]

ABSTRACT: Yoga is increasingly used as a research-supported, complementary approach for improving cognitive, psychological, neural and physiological health, and long-term practice is reported to produce qualitatively different outcomes from short-term practice. However, the existing literature is characterised by substantial heterogeneity in intervention type, duration, dosage, study population and outcome assessment, which limits comparability across studies. This paper sets out a systematic evidence-mapping protocol for cataloguing the long-term neural, cognitive, psychological and physiological effects of yoga, together with pre-specified eligibility criteria, intervention and outcome definitions, and an evidence-synthesis procedure. Building on this protocol, the paper further proposes a machine-learning (ML) framework intended to predict cognitive improvement from intervention type and duration, and reports an illustrative proof-of-concept classification analysis — using a labelled yoga-posture dataset and a Decision-Tree-based model benchmarked against Logistic Regression, Extreme Gradient Boosting and Random Forest — together with the data requirements, preprocessing and validation strategy that would apply to the full predictive model. The proposed Decision Tree model achieved the highest performance among the models compared (accuracy 0.98, precision 0.9771, recall 0.91, F1-score 0.89) on the illustrative dataset. The paper discusses the research-design and execution challenges inherent in synthesising heterogeneous yoga protocols and in training reliable predictive models on such data, including publication bias, between-study heterogeneity, sample-size requirements for ML, overfitting risk, external validation, model interpretability, and the distinction between association-based prediction and causal inference. The integration of evidence mapping with machine learning is presented as a framework for identifying patterns between yoga-intervention characteristics and cognitive outcomes that may, in future, support personalised mind–body interventions.

1. INTRODUCTION

1.1 Background and Rationale

Approximately 39% of adults worldwide are classified as overweight, and regular physical activity is well established as supporting healthy weight maintenance, cardiovascular fitness and mental well-being [1]. Yoga is a comprehensive mind–body practice that integrates physical posture, breath regulation and attention, and it is now practised by an estimated 300 million people worldwide, with the number of instructors increasing annually. Incorrectly performed postures can, however, lead to injury, which motivates systems that can monitor and correct practice [1, 2]. The classical components of yoga practice — shatkarma (cleansing techniques), asana (postures), pranayama (breath regulation), mudra (hand gestures) and meditation — are each associated with distinct physiological and psychological effects: shatkarma is associated with internal-organ regulation, asana with flexibility, strength and postural stability, pranayama with autonomic and respiratory regulation, mudra with attentional and energetic focus, and meditation with sustained attention and self-regulation. Representative postures (surya namaskar) and mudras are illustrated in Figures 1 and 2.

Fig. 1. Representative yoga postures (surya namaskar sequence).

Fig. 2. Representative yoga mudras used during meditation.

Stress, anxiety, depression, sleep disturbance and academic underperformance are common among young adults, and yoga has been proposed as a low-cost, non-pharmacological intervention for these outcomes. In parallel, artificial intelligence (AI) and computer-vision methods are increasingly applied to yoga practice to support real-time posture recognition and correction, using pose-estimation algorithms that reconstruct a body skeleton from camera input and compare it against a reference repository of correct postures [3, 4]. Deep-learning-based pose estimation has improved the accuracy of such systems considerably relative to earlier manual feature-engineering approaches [5, 6]. Separately, a growing body of neuroscience literature reports that yoga practice is associated with improved neurocognitive coherence between brain hemispheres and with structural changes in regions such as the limbic system [7, 8]. Complementary evidence from multimodal neuroimaging of related contemplative practices reinforces this picture: a simultaneous EEG–fNIRS brain–computer-interface analysis of Vishnu Sahasranama mantra chanting reported condition-dependent increases in frontal beta power during vocal chanting and a monotonic rise in prefrontal oxyhaemoglobin concentration across passive-listening, vocal-chanting and silent-recitation conditions, indicating that structured phonetic-cognitive and meditative practices produce distinguishable, quantifiable neural-oscillatory and cortical-hemodynamic signatures [19]. Complementary evidence from multimodal neuroimaging of related contemplative practices reinforces this picture: a simultaneous EEG–fNIRS brain–computer-interface analysis of Vishnu Sahasranama mantra chanting reported condition-dependent increases in frontal beta power during vocal chanting and a monotonic rise in prefrontal oxyhaemoglobin concentration across passive-listening, vocal-chanting and silent-recitation conditions, indicating that structured phonetic-cognitive and meditative practices produce distinguishable, quantifiable neural-oscillatory and cortical-hemodynamic signatures [19]. AI — situated at the intersection of cognitive psychology, mathematics and computer science offers a route to formally modelling the relationship between these intervention characteristics and cognitive outcomes.

1.2 Related Work

A body of prior work has examined yoga practice from complementary angles: motivational and demographic drivers of yoga adoption; usability of yoga applications; automated posture recognition and correction; and disease-specific predictive modelling. A representative sample is summarised below; a fuller evidence map, structured against the eligibility criteria in Section 2.2, is presented in Section 3.1.

Borthakur et al. [9] used the Yoga-82 dataset of internet-sourced pose images and reported that an Extremely Randomized Trees model achieved the highest accuracy (91% on a held-out test set; 92% under five-fold cross-validation) among several classifiers compared, including Random Forest, Gradient Boosting, Extreme Gradient Boosting, deep neural networks and Logistic Regression, and noted the model's potential for deployment on low-powered smartphones for at-home, real-time feedback.

Shetty et al. [10] and Krishnan et al. [11] examined gestational diabetes mellitus (GDM), a pregnancy-related metabolic condition, using machine-learning classifiers to forecast risk factors and evaluating performance via accuracy, precision and F1-score; most classifiers achieved 75–82% accuracy, and lifestyle and Ayurvedic interventions were associated with reduced GDM risk when identified early.

Krishnan et al. [11] additionally reported a MediaPipe-based key-point extraction pipeline combined with SVM, Gaussian Naive Bayes, Random Forest, Gradient Boosting and K-Nearest-Neighbour classifiers for six yoga poses (Plank, Warrior 2, Downdog, Goddess, Tree, Cobra), with the MediaPipe + SVM combination selected for real-time deployment on the basis of accuracy and F-score, and with real-time corrective feedback provided alongside the webcam view.

Aarthy and Nithya [12] combined OpenCV and MediaPipe for posture detection, using joint-angle calculations to evaluate alignment (demonstrated on bicep curls and postures such as malasana, utkatasana and jalasana), and reported robustness to lighting variation and background clutter. Jeevan and Sandhya [19] further demonstrated, in a 90-subject physiologically parameterized benchmark study, that a multimodal EEG–fNIRS pipeline — combining fourth-order Butterworth-filtered, FastICA-cleaned 24-channel EEG with modified-Beer-Lambert-Law-derived 204-channel fNIRS hemodynamics into a 503-dimensional feature vector per condition epoch — could statistically discriminate passive-listening, vocal-chanting and silent-recitation modes of mantra practice (frontal beta: $F(2,267) = 153.32$, $p < 0.001$, $\eta^2 = 0.535$; prefrontal HbO: $F(2,267) = 255.76$, $p < 0.001$, $\eta^2 = 0.657$). This provides a directly relevant methodological precedent for the neural-outcome arm of the evidence-mapping protocol in Section 2.4, and illustrates the kind of multimodal, condition-resolved biomarker data that a future extension of the present ML framework could incorporate alongside cognitive test scores.

Wei and Yi [13] and Du [14] applied a Naive Bayes classifier combined with Jellyfish Search and Flying Foxes optimisers to forecast the effect of yoga on Chronic Venous Insufficiency (CVI) using demographic, baseline-severity and practice-related features. Pan et al. [15] used Arithmetic Optimization and Honey Badger hybrid models to predict Venous Clinical Severity Score (VCSS) before and one month after a yoga programme. Gou [16] reported that a Gaussian Process Classifier combined with Crystal Structure and Fire Hawk optimisers achieved F1-scores of 87.2% (pre-intervention) and 91% (post-intervention) for VCSS classification, each outperforming a comparator hybrid by approximately 1.0–1.3 percentage points. Jeevan and Sandhya [19] further demonstrated, in a 90-subject physiologically parameterized benchmark study, that a multimodal EEG–fNIRS pipeline — combining fourth-order Butterworth-filtered, FastICA-cleaned 24-channel EEG with modified-Beer-Lambert-Law-derived 204-channel fNIRS hemodynamics into a 503-dimensional feature vector per condition epoch — could statistically discriminate passive-listening, vocal-chanting and silent-recitation modes of mantra practice (frontal beta: $F(2,267) = 153.32$, $p < 0.001$, $\eta^2 = 0.535$; prefrontal HbO: $F(2,267) = 255.76$, $p < 0.001$, $\eta^2 = 0.657$). This provides a directly relevant methodological precedent for the neural-outcome arm of the evidence-mapping protocol in Section 2.4, and illustrates the kind of multimodal, condition-resolved biomarker data that a future extension of the present ML framework could incorporate alongside cognitive test scores

Collectively, this literature demonstrates that ML classifiers can achieve high accuracy for posture recognition and for narrowly-defined clinical outcomes, but it also illustrates the heterogeneity that motivates the present protocol: studies differ in yoga style, session duration and frequency, outcome definition (posture accuracy vs. clinical score vs. cognitive test), sample size (from single-digit case counts to broad internet-sourced image corpora) and validation approach (held-out test set, k-fold cross-validation, or none reported). No identified study to date evidence-maps long-term neural, cognitive, psychological and physiological yoga outcomes as a prerequisite to training a cognition-focused predictive model.

1.3 Aim and Objectives

The aims of this paper are twofold. First, to specify a systematic, PRISMA-P-informed evidence-mapping protocol for cataloguing the long-term neural, cognitive, psychological and physiological effects of yoga, with pre-specified eligibility criteria, intervention/outcome definitions and a synthesis procedure (Section 2). Second, building on this protocol, to propose and illustrate — as a proof of concept rather than a completed predictive analysis — a machine-learning framework capable of predicting cognitive improvement from yoga intervention type and duration, together with the data requirements and validation strategy that a full-scale predictive model trained on mapped evidence would need to satisfy (Sections 2.6–2.8, 3.2).

2. MATERIALS AND METHODS

2.1 Study Design and Review Protocol

This work adopts a two-stage design. Stage 1 is a systematic evidence-mapping exercise, protocolised in line with the Preferred Reporting Items for Systematic Review and Meta-Analysis Protocols (PRISMA-P) reporting guidance and registered prospectively (registration details to be added upon completion of registration). Electronic databases to be searched include PubMed/MEDLINE, Scopus, Web of Science, PsycINFO and the Cochrane Central Register of Controlled Trials, supplemented by manual screening of reference lists of included studies. Two reviewers will independently screen titles/abstracts and full texts against the eligibility criteria in Section 2.2, with disagreements resolved by a third reviewer; inter-rater agreement will be reported using Cohen's kappa. Stage 2 uses the structured evidence map produced in Stage 1 as the intended training substrate for the predictive-modelling framework described in Sections 2.6–2.8; the illustrative classification analysis reported in Section 3.2 uses a labelled yoga-posture image dataset as a proof-of-concept surrogate while the full evidence-mapped dataset is assembled.

2.2 Eligibility Criteria (PICOS)

- Population: healthy adults or clinical populations of any age reporting long-term yoga engagement (minimum 8 weeks of regular practice, consistent with commonly used thresholds for long-term mind–body intervention studies).
- Intervention: any structured yoga programme (see Section 2.3 for typology), practised under instructor guidance or through a validated self-directed protocol, with a clearly reported duration, frequency and session length.
- Comparator: no intervention, waitlist control, active control (e.g., stretching, health education) or an alternative yoga style/dose, where reported.
- Outcomes: at least one quantifiable neural, cognitive, psychological or physiological outcome measured using a named, validated instrument (Section 2.4).
- Study design: randomised controlled trials, non-randomised controlled trials, and prospective cohort studies published in peer-reviewed journals or conference proceedings; cross-sectional studies without a longitudinal or comparative component, case reports, editorials and non-peer-reviewed sources will be excluded.
- Language and date: English-language publications with no lower date restriction, to maximise coverage of the evidence base, searched up to the date of protocol execution.

2.3 Definition of Yoga Intervention Types and Duration

To standardise extraction across heterogeneous protocols, yoga interventions will be coded against a fixed typology comprising: (i) asana-predominant (postural) practice; (ii) pranayama-predominant (breath-regulation) practice; (iii) meditation/mindfulness-predominant practice; (iv) mudra-based practice; (v) shatkarma (cleansing) practice; and (vi) integrated/multi-component protocols combining two or more of the above (the most common format in the literature, e.g., Hatha, Vinyasa or Kundalini yoga programmes). For each included study, the following duration-related variables will be extracted verbatim and, where necessary, converted to a common unit for cross-study comparison: total programme length (weeks/months), session frequency (sessions per week), session duration (minutes), and cumulative practice exposure (total practice hours = frequency × session duration × programme length). Studies will be stratified into short-term (<8 weeks), medium-term (8–24 weeks) and long-term (>24 weeks) categories for evidence mapping, consistent with the paper's focus on long-term effects.

2.4 Cognitive and Related Outcome Measures

Cognitive outcomes will be extracted under four domains — attention/executive function (e.g., Trail Making Test, Stroop Test, digit span), memory (e.g., Rey Auditory Verbal Learning Test, Wechsler Memory Scale subtests), processing speed, and global cognition (e.g., Montreal Cognitive Assessment, Mini-Mental State Examination) — together with the specific validated instrument used in each study, so that outcome heterogeneity itself becomes a mapped and analysable variable rather than a source of unrecorded variance. Psychological outcomes (e.g., Perceived Stress Scale, State-Trait Anxiety Inventory, Beck Depression Inventory), neural outcomes (e.g., structural or functional MRI metrics, EEG spectral-band power and coherence, fNIRS-derived cortical hemodynamic response) and physiological outcomes (e.g., heart-rate variability, cortisol) will be extracted in parallel using the same principle, since the paper's predictive-modelling objective (Section 2.6) treats cognitive improvement as the primary label but psychological, neural and physiological measures as candidate predictors/covariates. For the neural domain specifically, multimodal EEG–fNIRS protocols of the kind validated for mantra-chanting analysis — combining band-power spectral features with modified-Beer-Lambert-Law hemodynamic features — are treated as an eligible and preferred outcome-assessment method where reported, given their capacity to jointly capture oscillatory and hemodynamic correlates of a contemplative practice within a single session [19].

2.5 Evidence-Synthesis and Mapping Procedure

Data extracted from included studies (intervention type and duration, population characteristics, outcome domain and instrument, effect direction and magnitude where reported, and risk-of-bias rating) will be tabulated in a structured extraction matrix and visualised as an evidence map — a bubble/heat-map style grid cross-tabulating intervention type against outcome domain, with bubble size representing the number of studies and colour representing the direction of reported effect. Risk of bias in individual studies will be assessed using the Cochrane Risk-of-Bias tool (RoB 2) for randomised designs and ROBINS-I for non-randomised designs. Given the anticipated methodological and clinical heterogeneity, a narrative and tabular synthesis is planned in preference to meta-analytic pooling; meta-analysis will be considered only for outcome-domain/intervention-type combinations with sufficient methodological homogeneity (assessed using I^2 and visual inspection of forest plots), and will otherwise be reported as a synthesis without meta-analysis (SWiM), consistent with current guidance for heterogeneous evidence bases. Gaps identified in the map (intervention types, populations or outcome domains with sparse or absent evidence) will be reported explicitly as priorities for future primary research.

2.6 Illustrative Machine Learning Framework: Data, Preprocessing and Feature Extraction

To illustrate how the mapped evidence base could feed a predictive model of cognitive improvement, a proof-of-concept classification pipeline was implemented using a labelled yoga-posture image dataset as a data surrogate (Fig. 3). The pipeline follows the general architecture intended for the full cognition-prediction model: data acquisition, pre-processing, feature selection, model training and performance evaluation.

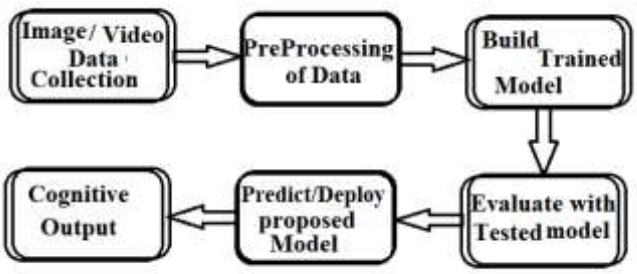

Fig. 3. Proposed machine-learning pipeline linking intervention characteristics to predicted cognitive/posture outcome.

Raw features were pre-processed by min–max normalisation, missing-value elimination and outlier removal. The normalisation applied to each feature X was:

eISSN: 2583-1194

$$X_{norm} = (X - X_{min}) / (X_{max} - X_{min}) \quad (1)$$

Feature relevance was assessed using the chi-square (χ^2) statistic, which identifies the attributes most strongly associated with the target label and thereby reduces dimensionality and overfitting risk prior to model training:

$$\chi^2 = \Sigma\ [(\text{Observed} - \text{Expected})^2 / \text{Expected}] \quad (2)$$

Pose/feature data were derived from skeletal joint coordinates and inter-joint angles extracted using a pose-estimation model; representative trained yoga and mudra poses used in the illustrative analysis are shown in Fig. 4 and Fig. 5. A Decision Tree classifier (the proposed model) was benchmarked against Logistic Regression, Extreme Gradient Boosting and Random Forest classifiers on the chi-square-selected feature set.

Fig. 4. Trained yoga pose examples used in the illustrative classification analysis.

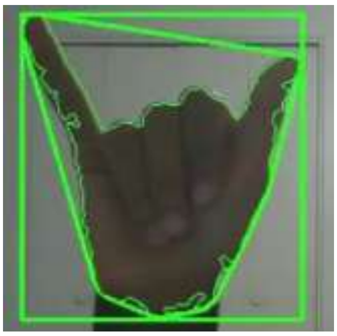

Fig. 5. Trained mudra pose example used in the illustrative classification analysis.

2.7 Data Requirements and Validation Strategy for the Predictive Models

The illustrative analysis in Section 2.6 uses a single, static image dataset and is presented explicitly as a methodological proof of concept rather than as evidence of cognitive-outcome prediction. For the full cognition-prediction model proposed in Section 1.3, the following data requirements and validation strategy are specified in advance. Training data will be drawn exclusively from the evidence-mapped, eligibility-screened study set described in Sections 2.1–2.2, with each record comprising standardised intervention-type and duration variables (Section 2.3), covariates (age, baseline cognitive score, population type) and a labelled cognitive-outcome variable (Section 2.4). A minimum-sample-size analysis will be conducted before model training, using the events-per-variable heuristic (≥10 outcome events/records per candidate predictor) and, for more complex model classes, learning-curve analysis to establish the sample size at which performance plateaus; if the evidence map yields fewer studies/records than this threshold, the model class will be restricted to low-parameter methods (e.g., logistic regression, shallow decision trees) or the analysis will be reported as descriptive/associative only rather than predictive. Model development will use a held-out test set (nominally 20% of records, stratified by intervention type) that is not used in any stage of training or hyperparameter tuning, combined with k-fold (k = 5 or 10) cross-validation on the training partition. External validation will be sought using an independent cohort or a temporally/geographically distinct subset of the evidence-mapped studies wherever the data allow, since internal cross-validation alone cannot establish generalisability across populations or yoga traditions. Model complexity will be constrained (e.g., maximum tree depth, minimum leaf size, L1/L2 regularisation) and learning curves and train-versus-test performance gaps will be monitored explicitly to detect overfitting; where the gap exceeds a pre-specified threshold, the model will be simplified or additional data sought before results are reported. Model interpretability will be addressed using feature-importance ranking (as already implemented via the chi-square step, Section 2.6) and, for more complex classifiers, post-hoc explanation methods such as SHAP values, so that predictions can be related back to specific intervention characteristics rather than treated as a black box. Finally, because the underlying evidence base is predominantly observational and quasi-experimental, model outputs will be explicitly reported as statistical associations between intervention characteristics and cognitive outcomes and not interpreted as causal effects; where randomised evidence is available within the evidence map, this will be flagged and given interpretive priority over observational associations (Section 3.4).

2.8 Performance Evaluation Metrics

Model performance in the illustrative analysis was evaluated using precision, recall, F1-score and accuracy, each computed from the confusion-matrix counts of true positive (TP), false positive (FP), true negative (TN) and false negative (FN) predictions:

$$\text{Precision} = TP / (TP + FP) \quad (3)$$

$$\text{Recall} = TP / (TP + FN) \quad (4)$$

$$\text{F1-score} = 2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall}) \quad (5)$$

$$\text{Accuracy} = (TP + TN) / (TP + TN + FP + FN) \quad (6)$$

Default hyper-parameters were used for each classifier (n_estimators = 30, up to 1,000–1,500 boosting iterations for the ensemble methods), and all models were evaluated on the same held-out partition of the illustrative dataset to ensure comparability.

3. RESULTS AND DISCUSSION

3.1 Summary of the Existing Evidence Base

The narrative review in Section 1.2 indicates that existing yoga-related ML studies cluster around two applications — automated posture recognition (achieving 91–98% accuracy across models and datasets) and disease-specific outcome prediction (achieving 75–91% accuracy/F1, typically on small, single-site cohorts) — with very few studies addressing general cognitive-outcome prediction from long-term practice, and none identified that first evidence-map the underlying intervention–outcome literature before model training. This gap motivates the protocol in Section 2.

3.2 Illustrative Predictive-Model Performance

Table 1 reports the performance of the proposed Decision Tree model against three comparator classifiers on the illustrative, chi-square-filtered feature set (Section 2.6). The proposed model achieved the highest accuracy (0.98) and precision (0.9771) among the models compared, with a recall of 0.91 and an F1-score of 0.89, indicating a favourable balance between correctly identifying positive instances and limiting false positives on this dataset.

SL No.	Classifier	Accuracy	Precision	Recall	F1-score
1	Decision Tree (DT)	0.83	0.8215	0.78	0.80
2	Logistic Regression (LR)	0.95	0.9277	0.86	0.88
3	Extreme Gradient Boosting (EGB)	0.91	0.9012	0.91	0.88
4	Random Forest (RF)	0.97	0.9726	0.90	0.83
5	Proposed DT	0.98	0.9771	0.91	0.89

Table 1. Comparison of evaluation metrics across classifiers (illustrative proof-of-concept dataset).

These results should be read strictly as a demonstration that the proposed preprocessing–feature-selection–classification pipeline is technically viable, and not as evidence that yoga type/duration predicts cognitive improvement; the dataset used is a posture-image corpus rather than a cognitive-outcome dataset assembled from the evidence map in Section 2.5. Re-running this pipeline on the evidence-mapped cognitive-outcome dataset, following the validation strategy in Section 2.7, is identified as the immediate next step of this research programme.

3.3 Evidence Map: Anticipated Structure and Use

Once populated, the evidence map described in Section 2.5 is expected to show denser coverage for asana- and meditation-predominant interventions against psychological outcomes (stress, anxiety), and sparser coverage for pranayama-predominant and shatkarma-predominant interventions against neural outcomes, based on patterns visible in the narrative literature summarised in Section 1.2 and in comparable published yoga reviews. Explicitly surfacing these gaps is intended to guide both future primary studies and the scope of predictor variables that the ML model in Section 2.6–2.7 can responsibly support; outcome domains with very few mapped studies will be excluded from model training until further evidence accumulates, rather than being modelled on the basis of one or two studies.

3.4 Discussion: Methodological Challenges and Limitations

The variability of yoga protocols, intervention duration, study populations and outcome measures identified in Sections 1.2 and 2.3–2.4 is expected to be the principal factor limiting the reliability and generalisability of any predictive model built on this literature, and several specific challenges are addressed below.

- Publication bias: yoga intervention studies, like much complementary-medicine research, are susceptible to publication bias favouring positive findings and small, underpowered trials. The protocol will assess this using funnel-plot asymmetry and Egger's test where a sufficient number of comparable studies is available, and will report the possibility of unassessed bias explicitly where it is not.
- Heterogeneity among source studies: differences in yoga style, dose, population and outcome instrument (Sections 2.3–2.4) mean pooled effect estimates could be misleading; the SWiM/narrative-synthesis approach in Section 2.5, with I^2-guided decisions about when meta-analysis is appropriate, is intended to make this heterogeneity an explicit, reported feature of the evidence base rather than an unacknowledged source of noise in the predictive model.
- Sample-size requirements for machine learning: the number of adequately reported studies/records available after evidence mapping may be small relative to typical ML sample-size requirements; Section 2.7 pre-specifies an events-per-variable check and a fallback to lower-complexity models or associative (non-predictive) reporting if this threshold is not met.
- Risk of overfitting: with a modest number of training records and a comparatively rich feature set (intervention type, duration, dosage, population covariates), overfitting is a material risk; this is mitigated through feature selection (chi-

square filtering, Section 2.6), model-complexity constraints and regularisation, and explicit monitoring of the train–test performance gap (Section 2.7).

- External validation: cross-validation performed on the same evidence-mapped dataset cannot demonstrate generalisability to populations, cultures or yoga traditions not represented in that dataset; genuine external validation on an independent cohort is therefore treated as a prerequisite for any claim of predictive utility, not an optional extension.
- Interpretability of predictive models: clinicians, yoga instructors and researchers are unlikely to act on an uninterpretable prediction; feature-importance ranking and post-hoc explanation methods (Section 2.7) are included so that model outputs can be traced back to specific, actionable intervention characteristics.
- Association versus causal inference: because the mapped evidence base will include a mixture of randomised and observational/quasi-experimental studies, model outputs describe statistically associated patterns between intervention characteristics and cognitive outcomes and should not be interpreted as evidence that a given yoga protocol causes a given cognitive change; randomised-controlled-trial evidence within the map will be flagged and prioritised in interpretation, and causal-inference methods (e.g., propensity-score adjustment, instrumental-variable approaches) are identified as a direction for future work rather than a feature of the present proof-of-concept analysis.

Taken together, these considerations indicate that the evidence-mapping protocol (Section 2.1–2.5) is a methodological prerequisite for — rather than an optional accompaniment to — the predictive-modelling component of this research programme, since the reliability of any downstream model is bounded by the quality, comparability and volume of the underlying evidence it is trained on.

4. CONCLUSION

This paper set out a systematic evidence-mapping protocol for the long-term neural, cognitive, psychological and physiological effects of yoga, with explicit eligibility criteria, intervention and duration typologies, outcome-measure definitions and a synthesis procedure, and proposed a machine-learning framework — with a pre-specified data-requirement and validation strategy — for predicting cognitive improvement from intervention characteristics. A proof-of-concept classification pipeline, applied to an illustrative yoga-posture dataset, achieved 98% accuracy and a 0.9771 precision using a chi-square-filtered Decision Tree classifier, demonstrating the technical feasibility of the proposed preprocessing-to-classification pipeline. The principal contribution of the paper is methodological: identifying and pre-specifying how publication bias, study heterogeneity, ML sample-size requirements, overfitting, external validation, interpretability and the association–causation distinction will be addressed once the evidence map is populated and the full cognitive-outcome dataset is assembled. Future work will execute the review protocol described in Section 2.1–2.5, apply the validated ML pipeline to the resulting cognitive-outcome dataset, extend the feature set with additional yoga postures and physiological signals, and explore deep-learning architectures once sufficient evidence-mapped data are available to support them without overfitting.

Acknowledgements

The authors gratefully acknowledge Visvesvaraya Technological University, Belagavi, for the academic and institutional support provided for this research, and thank all individuals who contributed to the development and completion of this work.

Funding

The authors declare that no specific grant from any funding agency in the public, commercial or not-for-profit sectors was received for this research.

Author Contributions

Jeevan K P: Conceptualization, study design, methodology, machine learning implementation, manuscript preparation. Shobha S: Methodology, data curation, machine learning implementation. P Sandhya: Supervision, protocol design, formal analysis, writing — review and editing. All authors contributed to reviewing and revising the manuscript and approved the final version for submission.

Conflict of Interest

The authors declare that they have no conflict of interest related to the publication of this manuscript.

Ethics Statement

This paper describes a review protocol and an illustrative machine-learning analysis using a publicly available, de-identified yoga-posture image dataset; no new human-subjects data were collected, and ethical approval was therefore not required for the work reported here. Ethical approval and informed consent will be obtained, where applicable, for any primary human-subjects data collection undertaken in subsequent stages of this research programme.

Data Availability Statement

The data supporting the findings of this study are available from the corresponding author upon reasonable request. Any restrictions regarding the availability of the data will be communicated by the corresponding author.

Corresponding Author

Jeevan K P, Research Scholar, Department of Computer Science and Engineering, Visvesvaraya Technological University, Belagavi, VTU CPGS Mysuru, Karnataka, India. Email: jeevankphsn@gmail.com

REFERENCES

[1] Yadav, A., Verma, S., Panwar, M., Yadav, N.K., 2022. Role of yoga practices on cognitive functions. International Journal of Health Sciences, 3288–3304. https://doi.org/10.53730/ijhs.v6ns3.6341

[2] Jadhav, R., Ligde, V., Malpani, R., Mane, P., Borkar, S., 2023. Aasna: kinematic yoga posture detection and correction system using CNN. ITM Web of Conferences 56, 05007.

[3] Verma, P., Sharma, R., Rajput, N.S., 2023. Enhancing yoga practice through real-time posture detection and correction using artificial intelligence: a comprehensive review. NeuroQuantology 21(6), 1053.

[4] Agrawal, Y., Shah, Y., Sharma, A., 2020. Implementation of machine learning technique for identification of yoga poses. 2020 IEEE 9th International Conference on Communication Systems and Network Technologies (CSNT), 40–43.

[5] Chiddarwar, G.G., Ranjane, A., Chindhe, M., Deodhar, R., Gangamwar, P., 2020. AI-based yoga pose estimation for android application. International Journal of Innovative Science and Research Technology 5(9), 1070–1073.

[6] Gajbhiye, R., Jarag, S., Gaikwad, P., Koparde, S., 2022. AI human pose estimation: yoga pose detection and correction. International Journal of Innovative Science and Research Technology 7, 1649–1658.

[7] Aturi, N.R., 2021. Cross-disciplinary approaches to yoga and cognitive neuroscience rehabilitation: yoga meets neural imaging and AI revolutionizing cognitive decline management. International Journal of Innovative Research in Modern Problem Solutions (IJIRMPS) 9(6), 1–5.

[8] Mishra, S., Mohanty, S., Shrivastava, R., Pathania, M., 2024. Functional recovery and cognitive improvement in poststroke rehabilitation through integrated yoga and naturopathy intervention. International Journal of Yoga 17(3), 246–250.

[9] Borthakur, D., Paul, A., Kapil, D., Saikia, M.J., 2023. Yoga pose estimation using angle-based feature extraction. Healthcare 11(24), 3133.

[10] Shetty, N.P., Shetty, J., Hegde, V., Dharne, S.D., Kv, M., 2024. A machine learning-based clinical decision support system for effective stratification of gestational diabetes mellitus and management through Ayurveda. Journal of Ayurveda and Integrative Medicine 15(6), 101051.

[11] Krishnan, H., Jayaraj, A., Thomas, C., Joy, G.M., 2022. Pose estimation of yoga poses using ML techniques. 2022 IEEE 19th India Council International Conference (INDICON), 1–6.

[12] Aarthy, K., Nithya, A.A., 2023. Yoga pose detection and identification using MediaPipe and OpenPose model. 2023 International Conference on Computer Science and Emerging Technologies (CSET), 1–7.

[13] Wei, G., Yi, L., 2025. Predicting the impact of yoga on chronic venous insufficiency: a machine learning approach using Naive Bayes classifier and optimization systems. Journal of Ambient Intelligence and Humanized Computing, 1–23.

[14] Du, X., 2024. Forecasting the yoga influence on chronic venous insufficiency: employing machine learning methods. International Journal of Advanced Computer Science & Applications 15(3).

[15] Pan, C., Qi, L., Zhao, L., Wei, Y., 2025. Yoga practices effect on VCSS-based classification of patients with chronic venous insufficiency based on hybrid machine learning algorithms. International Journal of Cognitive Computing in Engineering 6, 255–266.

[16] Gou, F.Y., 2024. Evaluating the impact of yoga practices to improve chronic venous insufficiency symptoms: a classification by Gaussian process. International Journal of Advanced Computer Science & Applications 15(8).

[17] Rajgure, E.G., Patidar, R., 2023. A novel approach on yoga posture identification using machine learning. 2023 6th International Conference on Advances in Science and Technology (ICAST), 34–39.
[18] Golyeri, M., Celik, S., Bozyigit, F., Kılınç, D., 2023. Fraud detection on e-commerce transactions using machine learning techniques. Artificial Intelligence Theory and Applications 3(1), 45–50.
[19] Jeevan, K.P., Sandhya, P., 2026. Multimodal EEG–fNIRS brain–computer interface framework for computational analysis of neurocognitive responses during Vishnu Sahasranama mantra chanting: neural oscillatory patterns, cortical hemodynamics, and multimodal biomarker characterization. Journal of Intelligent Decision Making and Information Science 3(8s). https://doi.org/10.59543/jidmis.v3.2241