ARTIFICIAL INTELLIGENCE IN CARDIOVASCULAR AUTONOMIC NEUROMODULATION EVIDENCE BOUNDARIES AND RESPONSIBLE CLINICAL TRANSLATION

Dr. Hakan Göçer [1], Gültekin Baş [2], Ahmet Barış Durukan [3], Halil İbrahim Yıldırımdemir [4]

[1] *Umit Hospital, Batıkent-Eskişehir/Türkiye, Türkiye, ORCID: 0000-0002-9644-9579, hgocer@gmail.com*
[2] *Umit Hospital, Batıkent-Eskişehir/Türkiye, Türkiye, ORCID: 0000-0001-6395-2947, gltknbs@gmail.com*
[3] *Liv Hospital Ankara, Türkiye, ORCID: 0000-0003-0566-0350, barisdurukan@yahoo.com*
[4] *Umit Hospital, Batıkent-Eskişehir/Türkiye, Türkiye, ORCID: 0000-0003-2116-5368, dryildirimdemir@gmail.com.tr*

Corresponding author: Dr. Hakan Göçer, hgocer@gmail.com

ABSTRACT

Cardiovascular autonomic neuromodulation brings together physiological treatment, behavioral support, and emerging computational technologies, but evidence for a treatment should not be mistaken for evidence that artificial intelligence improves it. This critical narrative review examines selected clinical trials and foundational neurotechnology studies alongside ethical and governance guidance. It distinguishes intervention efficacy, the incremental contribution of algorithmic adaptation, and the conditions for responsible implementation. Renal denervation has sham-controlled evidence for blood pressure reduction in selected populations; auricular stimulation and heart rate variability biofeedback have encouraging but more limited clinical evidence. Baroreflex activation requires qualified interpretation because a pivotal hypertension trial did not meet all prespecified primary endpoints. Cortical interfaces, decoded neurofeedback, and cardiac digital twins provide technological or conceptual context without establishing cardiovascular benefit from autonomous algorithmic control. No systematic search, pooled analysis, or formal certainty grading is claimed. The proposed translational framework requires comparison against competent nonadaptive care, clinically justified control limits, interpretable escalation, documented software changes, and meaningful patient participation. It also addresses mental privacy, equitable access, professional responsibility, and the organizational costs of monitoring. Artificial intelligence should therefore be evaluated as a potentially useful adjunct whose clinical and social value must be demonstrated for each intended use, rather than inferred from physiological plausibility or technological sophistication.

Keywords: : Artificial Intelligence, Autonomic Neuromodulation, Cardiovascular Care, Health Governance, Patient Autonomy

Received: 22 October 2026 **Accepted:** 2 November 2026 **Published:** 2 November 2026

Cite: Göçer, H., Baş, G., Durukan, A. B., & Yıldırımdemir, H. İ. (2026). Artificial Intelligence in Cardiovascular Autonomic Neuromodulation Evidence Boundaries and Responsible Clinical Translation. AI-Driven Journal of Interdisciplinary Social Research, 1(1), 70-81. https://aijisr.org/articles/AI-JISR-2026-0014

CC BY 4.0 — This article is published under the Creative Commons Attribution 4.0 International License.

1. Introduction

Autonomic regulation links cardiovascular physiology with emotion, behavior, and the conditions in which care is delivered. The neurovisceral integration model provides a conceptual account of these relationships, while physiological research describes reciprocal communication between the brain and cardiovascular system (Thayer & Lane, 2009; Silvani et al., 2016). These connections make autonomic neuromodulation an interdisciplinary subject. They do not imply that changing a neural signal necessarily improves a cardiovascular outcome.

Artificial intelligence (AI) may assist signal interpretation, identify candidate responders, or support adaptive treatment delivery. However, three questions must remain separate: whether a physiological intervention works, whether AI adds benefit over a competent alternative, and whether implementation preserves patient agency and accountable care. A positive answer to the first question cannot establish the other two. This distinction is particularly relevant when established device therapies are discussed alongside cortical interfaces or neurofeedback experiments.

This review develops an evidence-informed framework for clinical translation and healthcare governance. Its contribution is a structured separation of intervention evidence from claims about algorithmic enhancement, followed by proposed requirements for evaluating safety, autonomy, workload, and access. The discussion connects cardiovascular medicine with healthcare management, psychology, and the ethics of neurotechnology; it does not propose autonomous treatment for routine practice.

2. Review approach and limits of inference

This is a critical narrative review with conceptual analysis, rather than a systematic review or meta-analysis. Selected foundational studies, illustrative cardiovascular trials, and governance documents were checked against indexed bibliographic records, publisher pages, or official organizational sources during preparation on 22 October 2026. The clinical examples were chosen to represent different intervention mechanisms and levels of translation, including auricular stimulation, renal denervation, baroreflex activation, and behavioral biofeedback. Foundational neuroscience studies were retained to examine the limits of extrapolation.

The review does not claim an exhaustive search of subscription databases, a registered protocol, duplicate screening, or a reproducible count of eligible studies. No numerical risk-of-bias scores or formal Grading of Recommendations Assessment, Development and Evaluation assessments were performed. Consequently, the evidence descriptions in Table 1 are contextual appraisals of selected sources, not certainty ratings or a comprehensive ranking of therapies. No study-selection flow diagram is presented because verified screening counts are unavailable.

For each example, the analysis considers the population, comparator, endpoint, and whether the study actually tested an AI component. Empirical findings are distinguished from proposed applications and normative recommendations. Absence of an AI comparison in the cited studies is not treated as proof that no relevant study exists elsewhere. The literature remains selective, and the framework should be tested and revised as evidence develops.

Appendix A provides a proposed protocol for evaluating the incremental clinical value of AI in cardiovascular autonomic neuromodulation. It specifies eligibility, information sources, a reproducible PubMed strategy, screening, outcomes, risk-of-bias assessment, and synthesis. The protocol is informed by PRISMA-P (Moher et al., 2015); it has not been registered or executed and does not convert this narrative review into a completed systematic review.

3. Physiological targets and measurement boundaries

Cardiovascular neuromodulation includes interventions with different targets and risk profiles. Renal denervation (RDN) disrupts renal sympathetic nerves, whereas baroreflex activation therapy (BAT) electrically stimulates the carotid baroreflex pathway. Low-level tragus stimulation targets an auricular route associated with vagal afferent signaling. Heart rate variability (HRV) biofeedback uses behavioral training, often involving paced breathing. Their effects cannot be assumed to be interchangeable merely because each concerns autonomic regulation (Silvani et al., 2016).

HRV is a measurement domain, not a unitary measure of autonomic balance or treatment success. Interpretation depends on recording duration, rhythm, respiration, artifacts, medication, and the chosen index. Measurements during atrial fibrillation or frequent ectopy require particular caution because irregular intervals do not provide the same physiological information as normal sinus intervals. The relevant acquisition and analysis conditions should be prespecified rather than collapsed into a single wellness score (Task Force of the European Society of Cardiology and the North American Society of Pacing and Electrophysiology, 1996).

An algorithm that accurately estimates a physiological signal has therefore completed only one part of the translational task. Using that signal to select treatment requires an independently justified relationship between the measurement, the proposed action, and the clinical endpoint. A controller optimized to increase an HRV index could otherwise favor a measurable change without improving symptoms, blood pressure, arrhythmia burden, or patient experience. This is a proposed failure scenario, not an observed result of an implemented system.

4. Selected cardiovascular evidence

4.1 Auricular stimulation and atrial fibrillation

TREAT-AF randomized 53 patients with paroxysmal atrial fibrillation to active tragus or sham earlobe stimulation for one hour daily over six months. Median atrial fibrillation burden at six months was lower in the active group after baseline adjustment. This supports further evaluation of a defined stimulation protocol in a selected population; it does not establish routine effectiveness across all atrial fibrillation phenotypes or improvement in stroke or mortality (Stavrakis et al., 2020).

The distinction relevant to this review is that the trial evaluated stimulation, not an AI-driven controller. Its findings cannot be attributed to algorithmic personalization. Larger comparative studies would need to examine durability, adherence, stimulation tolerability, and patient-important outcomes before broad implementation. A computationally adjusted protocol would require its own comparator and safety assessment.

4.2 Renal denervation and hypertension

RDN has randomized sham-controlled evidence for blood pressure lowering. In RADIANCE II, 224 patients were randomized after withdrawal of antihypertensive medication; the baseline-adjusted between-group difference in daytime ambulatory systolic blood pressure at two months favored ultrasound denervation by 6.3 mmHg (Azizi et al., 2023). SPYRAL HTN-OFF MED Pivotal also evaluated catheter-based denervation against sham in the absence of antihypertensive medication (Böhm et al., 2020). These trials establish treatment-specific evidence rather than evidence that AI is needed for benefit.

The 2024 European Society of Cardiology hypertension guideline places RDN within selected, shared-decision contexts rather than as a universal substitute for medical therapy (McEvoy et al., 2024).

Proposed AI uses include response prediction and support for patient selection. Those applications require external validation and evidence that acting on a prediction improves decisions. Predicting blood pressure change accurately in a development dataset is insufficient to justify withholding an otherwise reasonable option or promoting an invasive procedure.

4.3 Baroreflex activation

BAT should not be assigned an unqualified high-certainty label on the basis of the selected literature. The Rheos pivotal trial randomized 265 implanted participants to immediate or delayed activation. It met the prespecified endpoints for sustained efficacy, BAT safety, and device safety, but did not meet the acute responder or procedural safety endpoints (Bisognano et al., 2011). The result supports a qualified discussion of blood pressure effects alongside the limitations of the trial and the risks of implantation.

Evidence for one device generation, population, or indication should not be generalized automatically to another. In particular, physiological plausibility does not establish an AI optimization strategy. Any proposed adaptive BAT protocol would need to demonstrate added value over appropriate conventional programming while retaining clinically justified limits.

4.4 Behavioral biofeedback

A randomized trial in 65 patients with uncomplicated hypertension compared behavioral neurocardiac training incorporating HRV biofeedback with an active relaxation control. It reported modest ambulatory blood pressure reductions and changes in measures of vagal heart rate modulation (Nolan et al., 2010). These findings apply to the studied hypertensive population and should not be extrapolated to heart failure without separate evidence. It does not establish prevention of cardiovascular events.

Biofeedback is especially relevant to implementation research because benefit depends on learning, participation, and sustained practice. An AI coach might adapt explanations or reminders, but the incremental contribution of those functions should be separated from breathing practice, clinician contact, and expectancy. Usability, language, accessibility, and the burden of repeated self-monitoring should be assessed alongside physiological outcomes.

Table 1 summarizes the selected evidence and the distinct questions posed by adding AI.

Table 1. Selected cardiovascular evidence and questions for AI enhancement

Strategy	Selected clinical evidence	Boundary of inference	AI question requiring evaluation
Auricular stimulation	TREAT-AF: sham-controlled study of 53 patients with paroxysmal atrial fibrillation.	Reduced measured burden under a defined protocol; no AI controller comparison.	Does adaptation improve burden or tolerability over fixed stimulation?
Renal denervation	RADIANCE II and SPYRAL HTN-OFF MED Pivotal: sham-controlled blood pressure trials.	Blood pressure benefit does not establish outcome benefit from AI selection.	Does validated prediction improve selection and decision quality?
Baroreflex activation	Rheos pivotal trial: 265 participants; not all coprimary endpoints met.	Device and indication matter; implantation risks remain relevant.	Does adaptive programming add benefit within safe limits?
HRV biofeedback	Nolan et al.: randomized active-control trial in 65 hypertensive patients.	Modest physiological effects; no demonstration of AI-specific benefit.	Does coaching improve sustained practice and patient-important outcomes?
Cortical interfaces and decoded neurofeedback	Engineering and neuroscience demonstrations outside the cited cardiovascular trials.	No direct cardiovascular efficacy inference from those demonstrations.	Is there a defined cardiovascular indication and a less invasive comparator?

Note. This table is not a systematic evidence rating; it is a contextual comparison of selected studies. Sources are discussed in Sections 4 and 5. HRV, heart rate variability; AI, artificial intelligence.

5. Neurotechnology concepts and limits of extrapolation

5.1 Cortical interfaces and decoded neurofeedback

The Neuralink platform report describes electrode and surgical engineering rather than a cardiovascular treatment trial (Musk & Neuralink, 2019). Decoded neurofeedback studies demonstrate perceptual learning or fear-related effects, not treatment of hypertension or arrhythmias; limited awareness of a reinforced neural representation should not be confused with absence of consent (Shibata et al., 2011; Koizumi et al., 2017). The cited cortical-interface and decoded-neurofeedback studies do not demonstrate cardiovascular benefit; the cited digital-twin perspective likewise does not establish therapeutic efficacy; therefore, cortical interfaces and decoded neurofeedback, together with cardiac digital twins discussed in Section 5.2, are outside the therapeutic scope of this review as cardiovascular treatments and are retained only to define the boundary of extrapolation. Covert semantic manipulation is not proposed as a cardiovascular treatment.

5.2 Digital twins and adaptive control

Cardiac digital twins offer a framework for integrating patient data with computational representations (Corral-Acero et al., 2020). Their use for autonomic treatment optimization remains a proposed application in this review. Model fit, physiological plausibility, and treatment utility are separate validation questions. A model may reproduce a recorded signal yet recommend an ineffective action when used outside the conditions on which it was fitted. Accordingly, the digital-twin example is retained only as computational context and does not provide evidence of cardiovascular benefit from AI-mediated neuromodulation.

Adaptive control should also be distinguished from AI. A fixed-rule controller can adjust stimulation using predefined thresholds without machine learning. Any learned controller should be compared with an appropriately designed rule-based alternative and expert programming. Without that

comparison, improvement might reflect more frequent measurement or better follow-up rather than a specific advantage of learning algorithms.

6. Governance and the organization of care

6.1 Meaningful consent and mental privacy

World Health Organization guidance and the neurorights literature provide complementary starting points for evaluating autonomy, privacy, and responsibility (World Health Organization, 2021; Ienca & Andorno, 2017). For cardiovascular applications, consent should identify what is measured, what the algorithm may change, who supervises it, and how the patient can pause participation. A signature alone does not establish understanding. Explanations should distinguish ordinary cardiovascular monitoring from neural recording and from inferences about psychological state.

A proposed minimum is purpose-limited collection with role-based access and a justified retention period. Raw signals and inferred labels may expose different information; a label such as stress or low adherence can affect how clinicians perceive a patient even when the underlying inference is uncertain. Patients should have a route to challenge consequential labels and discuss contextual explanations. These safeguards are design recommendations, not evidence that a particular product currently satisfies them.

6.2 Clinical responsibility and technical accountability

For an adaptive intervention, the clinical team should define the permitted action range before use. The system should identify missing or unreliable input, stop adjustment when safety conditions fail, and escalate to a responsible clinician. A manual override is meaningful only if staff can recognize the relevant failure, respond within the required time, and restore a safe state. The intervention-specific protocol should address plausible risks such as excessive slowing of heart rate, hypotension, discomfort, or delayed recognition of deterioration.

Accountability should be distributed explicitly across prescribing, monitoring, software maintenance, and incident investigation. A development log or model card should retain intended use, exclusions, training and validation populations where relevant, input requirements, known limitations, software versions, and update history. Clinical audit records should link the input, recommendation or action, safety checks, override, and resulting output. Exact reproducibility and access control should be tested rather than assumed.

The European Union AI Act provides a legal framework for AI systems, while U.S. Food and Drug Administration guidance on predetermined change control plans addresses planned modifications to AI-enabled device software (European Parliament & Council of the European Union, 2024; U.S. Food and Drug Administration, 2025). Applicability depends on the product, intended purpose, jurisdiction, and implementation. Neither framework should be cited as blanket approval of autonomous cardiovascular control. Proposed updates require defined verification, authorization, and rollback arrangements.

6.3 Equity workload and resource allocation

AI-mediated neuromodulation could widen access to support while also shifting costs and work to patients. Device ownership, connectivity, language, disability, and the ability to follow repeated instructions affect participation. A service should therefore evaluate assisted and nondigital options rather than treat smartphone access as an implicit eligibility criterion. For less-resourced languages,

community-informed translation with clinically competent human verification may be preferable to unreviewed automated instructions.

Monitoring also creates organizational obligations. Alerts must be triaged, failed transmissions investigated, and software changes reviewed. Evaluation should include clinician time, patient time, false alerts, treatment interruptions, training, maintenance, and translation costs. A decrease in face-to-face visits may conceal an increase in remote work. Economic assessment should compare the full care pathway with a credible alternative and examine who bears the costs.

Patient preference should be elicited after balanced information about benefit, uncertainty, privacy, and practical burden. High uptake does not establish free choice if alternatives are difficult to access. Clinicians should retain a documented reason for recommending an AI-assisted approach that goes beyond novelty or the availability of a device.

AI-mediated neuromodulation implementation checklist

This proposed checklist is for service planning before supervised evaluation; it is not a validated implementation instrument. The clinical lead should record an accountable owner, completion status, and supporting evidence for each item.

• Accessibility: assess disability, digital literacy, device availability, connectivity, and caregiver support; provide assisted and nondigital alternatives and test the patient's ability to pause participation.

• Language: provide clinically verified translations and accessible instructions; use teach-back to check understanding of benefit, uncertainty, alerts, and withdrawal options.

• Cost: document equipment, connectivity, patient time, staff monitoring, translation, training, maintenance, and replacement costs; specify funding and assess whether out-of-pocket costs exclude eligible patients.

• Staff training: assign prescribing, monitoring, technical, and escalation roles; verify competence using signal-failure, override, and emergency-response scenarios, with refresher training after material changes.

• Alert triage: define urgency categories, a named receiving team, coverage hours, response deadlines, backup routes, and emergency instructions; audit false alerts, missed alerts, response delays, and workload.

• Software update verification: record model and software versions; authorize updates only after regression, safety, and subgroup checks; verify compatibility, retain an audit trail, and test rollback before deployment.

Enrollment should wait until unresolved safety-critical items are addressed; access barriers and their mitigation should be documented and reviewed with patient representatives.

7. Research agenda

Phase 1 Technical validation

A staged programme should first specify one intended use, population, comparator, and clinical action. Synthetic or appropriately governed retrospective data can test signal failures and prohibited adjustments before patient-facing use. Technical validation should report missing-data handling, calibration where applicable, distribution shift, and subgroup performance. Passing these tests is necessary for further evaluation but does not establish clinical effectiveness. The following two proposed protocol templates operationalize control limits and interpretable escalation. They are

examples for adaptation to a single device and indication, not tested protocols or universal treatment thresholds.

Protocol template A Clinically justified control limits

• Intended use and design: evaluate advisory AI adjustment of auricular stimulation in adults with paroxysmal atrial fibrillation, first using governed retrospective signal replay and hardware-in-the-loop simulation. Compare a locked model with fixed-rule control and expert programming on identical test inputs; no patient stimulation occurs in this technical phase.

• Control specification: before testing, an independent clinical and engineering panel should document the permitted current, pulse width, frequency, duty cycle, and maximum rate of change for the chosen device. Define patient-specific heart-rate and blood-pressure stop criteria, symptoms requiring cessation, signal-quality limits, and missing-data timeouts using device specifications, relevant evidence, and clinical risk assessment. Record the rationale and exact numerical values in the approved protocol; generic values cannot be transferred across devices or indications.

• Stress tests and fallback: inject bradycardia, hypotension, ectopy, artifacts, delayed or missing inputs, transmission failure, and conflicting signals. Verify that an independent safety layer blocks prohibited commands, suspends adjustment when inputs are unreliable, and enters a prespecified device-appropriate safe state with manual override and clinician notification.

• Outcomes and progression: report prohibited-command attempts and escapes, stop-trigger sensitivity, detection-to-safe-state latency, false stops per monitored hour, and successful overrides with denominators and confidence intervals. Prespecify a maximum safe latency for each hazard. Any executed out-of-range command or failed safety-critical stop blocks progression pending correction and complete regression testing; absence of failures in a finite test set does not establish clinical safety.

Protocol template B Interpretable escalation

• Intended use and design: in a simulated monitoring service for the same auricular-stimulation use case, trained clinicians review randomized, matched scenarios with either the proposed explanatory alert or the existing alert display. Include both routine events and rare safety-critical failures; determine scenario and user numbers from the desired precision of the primary response measure.

• Alert specification: each alert should display the triggering observation and time, input quality and missingness, relevant clinical context, the proposed or withheld adjustment, the safety rule invoked, uncertainty where available, and the required action. Identify the responsible clinician, urgency, acknowledgement deadline, backup contact, and emergency route. Explanations should describe observable reasons for escalation rather than imply a proven causal interpretation of model features.

• Outcomes and progression: prespecify correct triage and appropriate action within the scenario-specific deadline as the primary human-factors outcome; also assess missed and false escalation, time to action, comprehension, overrides and reasons, automation-related errors, and staff workload. Independent adjudicators should review discordant decisions. Unrecognized safety-critical alerts or failed backup routing require redesign and retesting before live use.

• Reporting and live handoff for both templates: retain intended use, setting, participant and user eligibility, training, algorithm version, input handling, errors, overrides, modifications, and progression decisions. Use DECIDE-AI to structure the subsequent early live decision-support evaluation, including clinical safety and human factors (Vasey et al., 2022). Offline templates alone cannot demonstrate DECIDE-AI reporting compliance; interpretability is an additional proposed evaluation objective, not a specific mandated DECIDE-AI item.

Phase 2 Supervised clinical feasibility

A supervised feasibility study should then examine workflow, patient understanding, clinician workload, and the frequency and reasons for overrides. DECIDE-AI offers reporting guidance for early clinical evaluation where an AI component provides decision support; its relevance depends on the implemented system (Vasey et al., 2022). An autonomous controller may require additional device-specific evaluation. Protocols should prespecify clinically important failure categories, stopping criteria, and independent adjudication.

Phase 3 Comparative effectiveness

A subsequent randomized comparison should keep the physiological treatment and clinical contact as comparable as possible while varying algorithmic support. Fixed-rule control or expert programming is usually a more informative comparator than poorly monitored usual care. Primary outcomes should match the indication, such as ambulatory blood pressure or objectively measured arrhythmia burden; safety and patient-important outcomes should be measured separately. Sample size and follow-up must reflect the chosen endpoint, and a small pilot should not be described as establishing rare-event safety.

Phase 4 Implementation and external validation

Implementation studies should assess whether any clinical gain persists after accounting for costs, accessibility, trust, and staff capacity. Prespecified subgroup analyses should examine differential error and treatment burden, supported by qualitative interviews where useful. Long-term cardiovascular events require adequate power and follow-up; neither improved HRV nor patient satisfaction alone can substitute for those outcomes. These studies would test the proposed framework rather than validate it by assumption.

External validation plans for the AI questions in Table 1

For every proposed AI application, lock the model and decision thresholds before evaluation in geographically separate centres and a later time period, with no patient overlap or model tuning on the external test set. Prespecify eligibility, sample size based on precision or the clinically important effect, missing-data handling, subgroup analyses, and adjudicated outcomes. Report uncertainty and site-specific performance; any subsequent recalibration requires a separate evaluation dataset.

• Auricular stimulation: first test recommendation safety and signal robustness using independent recordings and device settings; then compare supervised adaptive with fixed stimulation while holding contact time constant. Assess objectively measured atrial fibrillation burden, tolerability, adherence, and interruptions.

• Renal denervation: validate a response-prediction model across centres with different baseline blood pressure, medication regimens, and procedural characteristics, reporting calibration and prediction error against follow-up ambulatory blood pressure. A subsequent prospective decision-impact comparison should assess selection quality, treatment uptake, inappropriate exclusion, and patient outcomes. Predicting observed response alone does not establish an individual treatment effect.

• Baroreflex activation: validate only within the specified device generation and indication, using independent data and bench tests before supervised comparison with expert programming. Assess blood pressure response, symptoms, out-of-range commands, overrides, and device-related adverse events without extrapolating across indications.

• HRV biofeedback: evaluate the coaching system in independent clinics and language groups against standardized non-AI coaching with comparable contact and breathing instruction. Assess sustained practice, ambulatory blood pressure, patient-important outcomes, accessibility, and monitoring burden; HRV acquisition conditions should remain standardized.

• Cortical interfaces and decoded neurofeedback: no cardiovascular deployment validation is proposed at present. A justified indication, plausible mechanism, appropriate consent, and comparison with less invasive alternatives must precede any clinical protocol. Independent replication of a relevant cardiovascular effect would be required before comparative therapeutic evaluation.

8. Limitations

This review uses selected sources and may omit relevant studies, including recent developments. It does not estimate pooled effects, compare all available devices, or provide a formal guideline recommendation. Clinical populations and endpoints differ substantially, which limits cross-intervention comparison. The governance framework is an analytical proposal and has not been evaluated in an operating service. Economic effects, patient preferences, and the incremental benefit of AI remain questions for prospective study. The added checklist, protocol templates, and external validation plans are prospective proposals and have not been empirically validated. Appendix A provides a proposed systematic review protocol for future evidence synthesis. The search has not been expanded or executed systematically in the present review, and the protocol has not been registered. Its prospective implementation would allow the completeness and certainty of the evidence to be evaluated without changing the methodological designation of the present article.

9. Conclusion

Cardiovascular neuromodulation provides several distinct settings in which AI might assist measurement, selection, or adaptation. Existing intervention evidence cannot establish that those computational additions improve care. Responsible translation requires an indication-specific comparison with competent non-AI alternatives and evaluation of safety, patient agency, workload, and equitable access. Cortical interfaces, decoded neurofeedback, and digital twins should remain clearly separated from demonstrated cardiovascular treatment effects until the relevant clinical evidence is available.

Ethics approval

Not applicable. This narrative review and conceptual analysis involved no original human or animal research and no identifiable patient data.

Informed consent

Not applicable.

Funding

No specific funding was received for preparation of this article.

Conflict of interest

The author declares no conflicts of interest.

Data availability

No original datasets were generated or analyzed. Sources used for this narrative review are listed in the references.

Use of AI tools

ChatGPT (OpenAI) was used to assist language editing, restructuring, reference verification support, and document formatting. No patient data were processed and no synthetic clinical results were generated. Responsibility for the scientific interpretation, source verification, and approval of the submitted text rests with the human author.

Author contributions

CRediT roles: Conceptualization, Investigation, Writing – original draft, Writing – review and editing. Author identification is supplied separately through the submission system to preserve blinded review.

References

Azizi, M., Saxena, M., Wang, Y., Jenkins, J. S., Devireddy, C., Rader, F., Fisher, N. D. L., Schmieder, R. E., Mahfoud, F., Lindsey, J., Sanghvi, K., Todoran, T. M., Pacella, J., Flack, J., Daemen, J., Sharp, A. S. P., Lurz, P., Bloch, M. J., Weber, M. A., … RADIANCE II Investigators and Collaborators. (2023). Endovascular ultrasound renal denervation to treat hypertension: The RADIANCE II randomized clinical trial. JAMA, 329(8), 651-661. https://doi.org/10.1001/jama.2023.0713

Bisognano, J. D., Bakris, G., Nadim, M. K., Sanchez, L., Kroon, A. A., Schafer, J., de Leeuw, P. W., & Sica, D. A. (2011). Baroreflex activation therapy lowers blood pressure in patients with resistant hypertension: Results from the double-blind, randomized, placebo-controlled Rheos pivotal trial. Journal of the American College of Cardiology, 58(7), 765-773. https://doi.org/10.1016/j.jacc.2011.06.008

Böhm, M., Kario, K., Kandzari, D. E., Mahfoud, F., Weber, M. A., Schmieder, R. E., Tsioufis, K., Pocock, S., Konstantinidis, D., Choi, J. W., East, C., Lee, D. P., Ma, A., Ewen, S., Cohen, D. L., Wilensky, R., Devireddy, C. M., Lea, J., Schmid, A., … SPYRAL HTN-OFF MED Pivotal Investigators. (2020). Efficacy of catheter-based renal denervation in the absence of antihypertensive medications (SPYRAL HTN-OFF MED Pivotal): A multicentre, randomised, sham-controlled trial. The Lancet, 395(10234), 1444-1451. https://doi.org/10.1016/s0140-6736(20)30554-7

Corral-Acero, J., Margara, F., Marciniak, M., Rodero, C., Loncaric, F., Feng, Y., Gilbert, A., Fernandes, J. F., Bukhari, H. A., Wajdan, A., Martinez, M. V., Santos, M. S., Shamohammdi, M., Luo, H., Westphal, P., Leeson, P., DiAchille, P., Gurev, V., Mayr, M., … Lamata, P. (2020). The 'Digital Twin' to enable the vision of precision cardiology. European Heart Journal, 41(48), 4556-4564. https://doi.org/10.1093/eurheartj/ehaa159

European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

Ienca, M., & Andorno, R. (2017). Towards new human rights in the age of neuroscience and neurotechnology. Life Sciences, Society and Policy, 13(1), 5. https://doi.org/10.1186/s40504-017-0050-1

Koizumi, A., Amano, K., Cortese, A., Shibata, K., Yoshida, W., Seymour, B., Kawato, M., & Lau, H. (2017). Fear reduction without fear through reinforcement of neural activity that bypasses conscious exposure. Nature Human Behaviour, 1, 0006. https://doi.org/10.1038/s41562-016-0006

McEvoy, J. W., McCarthy, C. P., Bruno, R. M., Brouwers, S., Canavan, M. D., Ceconi, C., Christodorescu, R. M., Daskalopoulou, S. S., Ferro, C. J., Gerdts, E., Hanssen, H., Harris, J., Lauder, L., McManus, R. J., Molloy, G. J., Rahimi, K., Regitz-Zagrosek, V., Rossi, G. P., Sandset, E. C., … Khamidullaeva, G. A. (2024). 2024 ESC Guidelines for the management of elevated blood pressure and hypertension. European Heart Journal, 45(38), 3912-4018. https://doi.org/10.1093/eurheartj/ehae178

Moher, D., Shamseer, L., Clarke, M., Ghersi, D., Liberati, A., Petticrew, M., Shekelle, P., Stewart, L. A., & PRISMA-P Group. (2015). Preferred reporting items for systematic review and meta-analysis protocols (PRISMA-P) 2015 statement. Systematic Reviews, 4, 1. https://doi.org/10.1186/2046-4053-4-1

Musk, E., & Neuralink. (2019). An integrated brain-machine interface platform with thousands of channels. Journal of Medical Internet Research, 21(10), e16194. https://doi.org/10.2196/16194

Nolan, R. P., Floras, J. S., Harvey, P. J., Kamath, M. V., Picton, P. E., Chessex, C., Hiscock, N., Powell, J., Catt, M., Hendrickx, H., Talbot, D., & Chen, M. H. (2010). Behavioral neurocardiac training in hypertension: A randomized, controlled trial. Hypertension, 55(4), 1033-1039. https://doi.org/10.1161/hypertensionaha.109.146233

Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., … Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. BMJ, 372, n71. https://doi.org/10.1136/bmj.n71

Shibata, K., Watanabe, T., Sasaki, Y., & Kawato, M. (2011). Perceptual learning incepted by decoded fMRI neurofeedback without stimulus presentation. Science, 334(6061), 1413-1415. https://doi.org/10.1126/science.1212003

Silvani, A., Calandra-Buonaura, G., Dampney, R. A., & Cortelli, P. (2016). Brain-heart interactions: Physiology and clinical implications. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 374(2067), 20150181. https://doi.org/10.1098/rsta.2015.0181

Stavrakis, S., Stoner, J. A., Humphrey, M. B., Morris, L., Filiberti, A., Reynolds, J. C., Elkholey, K., Javed, I., Twidale, N., Riha, P., Varahan, S., Scherlag, B. J., Jackman, W. M., Dasari, T. W., & Po, S. S. (2020). TREAT AF (Transcutaneous Electrical Vagus Nerve Stimulation to Suppress Atrial Fibrillation): A randomized clinical trial. JACC: Clinical Electrophysiology, 6(3), 282-291. https://doi.org/10.1016/j.jacep.2019.11.008

Task Force of the European Society of Cardiology and the North American Society of Pacing and Electrophysiology. (1996). Heart rate variability: Standards of measurement, physiological interpretation and clinical use. Circulation, 93(5), 1043–1065. https://doi.org/10.1161/01.CIR.93.5.1043

Thayer, J. F., & Lane, R. D. (2009). Claude Bernard and the heart-brain connection: Further elaboration of a model of neurovisceral integration. Neuroscience & Biobehavioral Reviews, 33(2), 81-88. https://doi.org/10.1016/j.neubiorev.2008.08.004

U.S. Food and Drug Administration. (2025). Marketing submission recommendations for a predetermined change control plan for artificial intelligence-enabled device software functions: Guidance for industry and Food and Drug Administration staff. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/marketing-submission-recommendations-predetermined-change-control-plan-artificial-intelligence

Vasey, B., Nagendran, M., Campbell, B., Clifton, D. A., Collins, G. S., Denaxas, S., Denniston, A. K., Faes, L., Geerts, B., Ibrahim, M., Liu, X., Mateen, B. A., Mathur, P., McCradden, M. D., Morgan, L., Ordish, J., Rogers, C., Saria, S., Ting, D. S. W., … DECIDE-AI expert group. (2022). Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. Nature Medicine, 28(5), 924-933. https://doi.org/10.1038/s41591-022-01772-9

World Health Organization. (2021). Ethics and governance of artificial intelligence for health: WHO guidance. https://www.who.int/publications/i/item/9789240029200