

Review

Not peer-reviewed version

---

# Artificial Intelligence Based Risk Stratification in Obesity Care: From Diagnosis to Personalised Treatment Pathways

---

[Simona Wójcik](#)\*, Monika Tomaszewska, Anna Rulkiewicz

Posted Date: 11 March 2026

doi: 10.20944/preprints202603.0855.v1

Keywords: AI; machine learning; risk stratification; phenotyping; digital therapeutics; behavioural coaching; electronic health records; wearable devices; genomics; clinical decision support



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Review

# Artificial Intelligence Based Risk Stratification in Obesity Care: From Diagnosis to Personalised Treatment Pathways

Simona Wójcik <sup>1,2,\*</sup>, Monika Tomaszewska <sup>2</sup> and Anna Rulkiewicz <sup>1</sup>

<sup>1</sup> LUX MED, ul. Szturmowa 2, 02-678 Warsaw, Poland

<sup>2</sup> Wyższa Szkoła Nauk Medycznych LUX MED, ul. Szturmowa 2, 02-678 Warsaw, Poland

\* Correspondence: simona.wojcik@luxmed.pl; +48-606121509

## Abstract

**Background/Objectives:** Obesity is a chronic, relapsing disease with a widening gap between clinical need and the availability of specialist care. Artificial intelligence (AI) may enable earlier risk detection, more precise phenotyping, and scalable behavioural support across obesity treatment pathways. This narrative review synthesises contemporary AI applications across the obesity care continuum and evaluates their translational readiness. **Methods:** A targeted search of PubMed/MEDLINE and Google Scholar (January 2024–January 2026) was conducted, complemented by citation chaining. Evidence was synthesised across four domains: (1) risk prediction and screening, (2) environmental and behavioural determinants, (3) multimodal phenotyping and precision stratification, and (4) AI-enabled lifestyle interventions and behavioural coaching (AIBC). **Results:** EHR-based models demonstrate clinically useful discrimination for early risk identification. Multimodal approaches refine stratification beyond BMI-centric classification. AIBC platforms show emerging evidence of clinically meaningful weight loss, including non-inferiority to human coaching, but long-term effectiveness, generalisability, and equity remain insufficiently established. **Conclusions:** AI is positioned to become a core enabler of personalised obesity pathways. Safe translation requires external validation, bias auditing, transparent reporting, human oversight, and post-deployment surveillance aligned with clinical guidelines and regulatory expectations.

**Keywords:** AI; machine learning; risk stratification; phenotyping; digital therapeutics; behavioural coaching; electronic health records; wearable devices; genomics; clinical decision support

---

## 1. Introduction

### 1.1. Background and Rationale

Obesity has emerged as one of the defining health challenges of the 21st century. Recent global estimates suggest that approximately one in eight individuals are living with obesity, with prevalence rising across virtually all age groups and regions. The economic burden is equally alarming: obesity and its related complications are estimated to account for health-care and productivity costs comparable to several percent of global GDP. The impact is not evenly distributed. A growing majority of people with obesity now live in low- and middle-income countries, where rapid nutrition transition, urbanization, and constrained health-system capacity intersect. Against this backdrop, traditional models of obesity care – episodic visits, limited contact time, and fragmented services – are poorly suited to managing a complex, chronic and relapsing condition that requires continuous monitoring, behavioural support, and timely adjustment of therapy.

At the same time, there is a profound mismatch between the scale of need and the availability of specialized obesity care. Multidisciplinary services capable of delivering evidence-based lifestyle interventions, pharmacotherapy, and metabolic and bariatric surgery remain scarce and unevenly

distributed, even in high-income settings. Long waiting times, geographic barriers, and workforce shortages further limit access. The result is a persistent treatment gap: large numbers of individuals with obesity either receive no structured care at all or only intermittent, low-intensity interventions that are unlikely to produce durable weight loss or meaningful improvements in obesity-related complications. There is an urgent need for scalable, cost-effective approaches that can extend the reach of specialist expertise, personalize treatment intensity to individual risk and response, and provide continuity of support between in-person encounters.

Artificial intelligence (AI) has rapidly advanced from a research topic to a practical component of contemporary health care, offering precisely the kinds of capabilities that could help address these gaps. In diagnostic disciplines such as radiology, pathology, dermatology and cardiology, machine learning and deep learning systems now routinely achieve performance comparable to human experts in image interpretation and pattern recognition. Beyond diagnostics, AI is increasingly embedded in electronic health records and hospital information systems, where it supports tasks such as clinical coding, triage, risk prediction, and resource allocation. Generative AI models, including large language models, have begun to transform clinical documentation workflows by drafting encounter notes, discharge summaries, and referral letters, reducing administrative burden and potentially alleviating clinician burnout. The breadth of these emerging capabilities—from medical note generation and patient triage to consultation support and preventive programmeme design—has been outlined in early appraisals of large language models in clinical settings, highlighting both the transformative potential and the persistent limitations of current systems, including the inability to replace direct physician–patient interaction [1].

Taken together, these advances reflect a qualitative shift in AI's role in medicine: from experimental tool to embedded infrastructure. In obesity care specifically, where sustained behaviour change, early risk identification, and responsive treatment adaptation are central but chronically under-resourced, AI offers capabilities that are directly relevant—not merely aspirational. Integration of multimodal data from electronic records and wearable devices creates new opportunities for continuity of support that episodic clinic visits cannot provide

### *1.2. Aim and Scope*

The aim of this narrative review is to synthesize current applications of artificial intelligence across the obesity care continuum and to assess their translational readiness for routine clinical use. Specifically, we examine four complementary domains: (1) early risk prediction and screening using electronic health records and population-based cohorts; (2) modelling of environmental and behavioural determinants of obesity, including image-based characterization of the built environment; (3) multidimensional phenotyping and precision risk stratification using multimodal data such as genomics, imaging and other omics profiles; and (4) AI-enabled lifestyle interventions and AI-enabled behavioural coaching (AIBC) as scalable mechanisms for continuous, personalised support between in-person encounters. In parallel, we discuss implementation and governance considerations, focusing on patient safety, bias and equity, transparency, human oversight and post-deployment monitoring, and we outline practical requirements for validating and integrating AI tools into guideline-driven obesity management and multidisciplinary care pathways.

### *1.3. Definitions and Functional Taxonomy of AI in Obesity Care*

To reduce conceptual ambiguity and to maintain consistency across heterogeneous evidence, we organize AI applications by clinical function within obesity pathways (prediction, stratification, decision support, behavioural delivery, and generative interfaces).

In obesity research and clinical practice, “artificial intelligence” refers to a heterogeneous set of computational approaches that differ in clinical purpose, data modalities, level of autonomy and position within care pathways. To improve conceptual clarity, this review adopts a functional taxonomy that organizes AI applications by what they do within obesity care.

First, predictive ML/DL models aim to forecast risk and clinically relevant outcomes. In obesity, this includes prediction of incident obesity within predefined windows, risk of cardiometabolic complications, likelihood of response to lifestyle or pharmacological treatment, and longer-term weight trajectories or relapse. These models are typically trained on electronic health records, cohort data and wearable-derived signals, and they serve as “upstream” components that support early detection and triage.

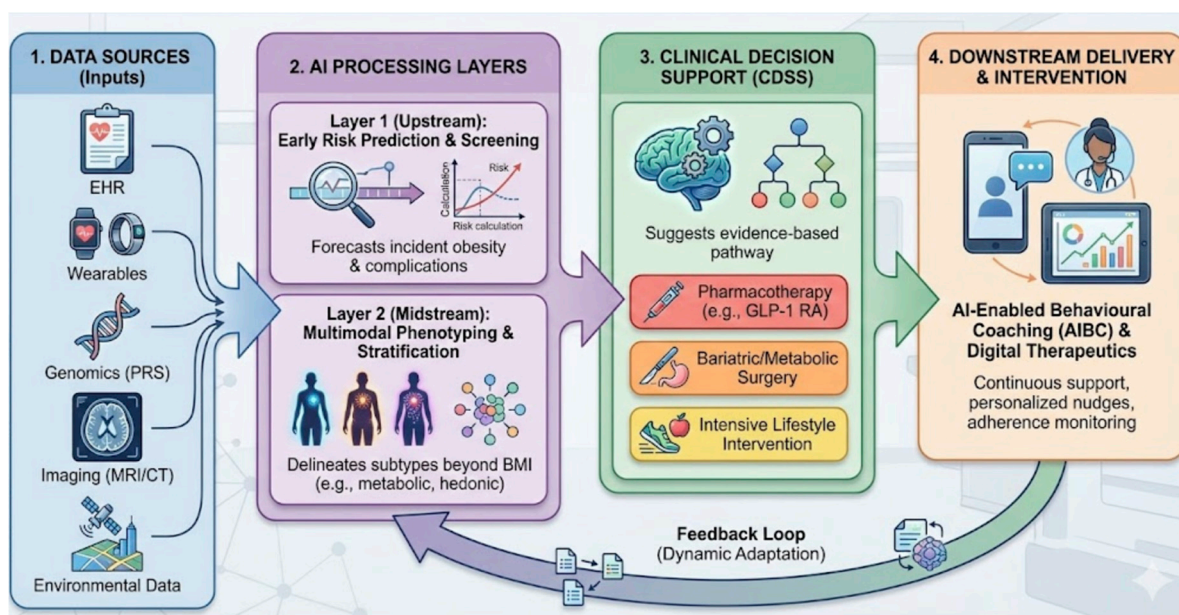
Second, AI supports risk stratification and multidimensional phenotyping beyond BMI-centric classification. Phenotyping models integrate multimodal data, including imaging-derived measures of visceral and ectopic adiposity (MRI, CT, DXA), metabolic biomarkers and omics profiles (e.g., metabolomics, microbiome), and genetic information (single nucleotide polymorphisms and polygenic risk scores). The goal is to delineate biologically and clinically meaningful subtypes of obesity that may differ in pathophysiology, prognosis and treatment responsiveness.

Third, clinical decision support systems (CDSS) and pathway-level orchestration tools embed predictive and phenotyping outputs into structured workflows. Examples include procedure selection support (pharmacotherapy versus metabolic and bariatric surgery), escalation or de-escalation of intervention intensity, monitoring strategies and multidisciplinary care coordination. In this view, AI is not merely a model but part of an operational architecture that mediates how evidence-informed decisions are executed over time.

Fourth, AI-enabled lifestyle interventions and AI-enabled behavioural coaching (AIBC) operationalize continuous, personalised behavioural support. AIBC systems may combine conversational agents, just-in-time adaptive interventions (JITAI), hybrid “digital twin” components and engagement strategies such as gamification [2]. Their distinguishing feature is dynamic tailoring of content, frequency and tone based on multimodal inputs from wearables, self-reports and contextual signals, enabling a persistent “background layer” of support across treatment modalities.

Fifth, generative AI, particularly large language models, functions as a human-facing interface between AI infrastructure and clinicians and patients. Typical use cases include documentation drafting, generation of patient education materials aligned with literacy and language needs, and interactive communication support. However, generative systems introduce distinct risks, notably hallucinations and guideline discordance, which necessitate explicit safeguards and human oversight. Empirical evaluations of LLM performance on standardised medical examinations have demonstrated that these systems can achieve passing-level accuracy across clinical domains, yet exhibit systematic weaknesses in questions requiring nuanced clinical reasoning and image interpretation, underscoring the need for domain-specific validation before deployment[3]

Finally, from an implementation perspective it is useful to distinguish clinician-facing tools (supporting professional decision-making) from patient-facing tools (delivering recommendations directly to patients), as these categories differ materially in acceptable risk thresholds, regulatory expectations and requirements for human-in-the-loop oversight. This taxonomy provides the organizing frame for the evidence synthesis that follows (Figure 1).



**Figure 1.** Functional architecture of AI-enabled obesity care pathways. The framework illustrates four operational layers: (1) data inputs (EHR, wearables, genomics, imaging, and environmental data); (2) AI processing layers comprising upstream early risk prediction and screening and midstream multimodal phenotyping and stratification; (3) clinical decision support (CDSS) generating evidence-based pathway recommendations across pharmacotherapy, bariatric/metabolic surgery, and intensive lifestyle intervention; and (4) downstream delivery through AI-enabled behavioural coaching (AIBC) and digital therapeutics, with a dynamic adaptation feedback loop linking monitoring outputs back to upstream models. Abbreviations: EHR, electronic health record; PRS, polygenic risk score; MRI, magnetic resonance imaging; CT, computed tomography; GLP-1 RA, glucagon-like peptide-1 receptor agonist; AIBC, AI-enabled behavioural coaching; CDSS, clinical decision support system.

#### 1.4. Roadmap of the Manuscript

In the sections that follow, we first summarize upstream AI for risk prediction and screening, then review environmental and behavioural modelling, and subsequently discuss multimodal phenotyping and precision stratification. We then evaluate downstream delivery tools, including AIBC and digital therapeutics, and conclude with implementation requirements for safety, equity, governance and clinical integration.

## 2. Methods

This manuscript is a narrative review. The narrative review format was selected to synthesize rapidly evolving and methodologically heterogeneous evidence on artificial intelligence (AI) applications in obesity care, spanning predictive modelling, multidimensional phenotyping, AI-enabled behavioural coaching, and clinical governance. Given the breadth of data modalities (EHR, imaging, wearables, genomics, and digital therapeutics) and the variability in study designs, populations and outcomes, a narrative approach was deemed the most appropriate to integrate findings across domains and to identify convergent themes, translational gaps and implementation considerations.

A targeted literature search was conducted in PubMed/MEDLINE and Google Scholar. The search was deliberately restricted to publications from January 2024 to January 2026. This narrow temporal window was selected for three reasons: first, the AI landscape in obesity care has undergone a qualitative shift in this period, marked by the first regulatory clearances of AI-driven digital therapeutics for weight management, the publication of the WHO guideline on GLP-1 RA incorporating behavioural adjuncts, and the emergence of landmark RCTs comparing AI-enabled coaching with human-delivered programmes. Second, the breadth of AI applications in clinical

settings—including documentation, triage, and patient communication—has been outlined in early appraisals of large language models in healthcare [1,3], while AI specifically in obesity has been addressed in recent comprehensive reviews [4,5]; restricting our search window avoids redundant coverage while ensuring focus on the translational frontier. Third, the rapid evolution of generative AI—including the clinical deployment of large language models for patient education and documentation—has introduced an entirely new category of tools that did not exist in earlier review periods. The search was last updated on 05.02.2026 to ensure inclusion of the most recent evidence available at the time of manuscript preparation. Key search terms were combined using Boolean operators and included: “obesity” OR “adiposity” OR “bariatric” AND “artificial intelligence” OR “machine learning” OR “deep learning” OR “large language model” OR “generative AI” OR “digital therapeutic” OR “behavioural coaching” OR “JITAI” OR “wearables” OR “EHR” OR “electronic health record” OR “phenotyping” OR “polygenic risk score” OR “microbiome” OR “metabolomics” OR “MRI” OR “CT” OR “DXA”. Reference lists of highly relevant reviews and consensus statements were additionally screened to identify further eligible studies (snowballing).

We included: (i) peer-reviewed review articles, consensus statements and position papers addressing AI in obesity, metabolic and bariatric surgery, or obesity-related cardiometabolic risk; (ii) original studies using ML/DL to predict obesity risk, treatment response or clinically relevant outcomes; (iii) studies on AI-enabled behavioural coaching or AI-driven digital lifestyle interventions, including randomised controlled trials when available; and (iv) studies integrating multimodal data for obesity phenotyping or risk stratification (e.g., genomics, imaging, metabolomics, microbiome).

We excluded: (i) purely technical papers without clinical outcomes or clearly defined health-related endpoints; (ii) non-obesity populations unless the primary aim was directly transferable to obesity management (e.g., diabetes prevention programmes with a substantial overweight/obesity subgroup); (iii) editorials without substantive synthesis; and (iv) non-English publications.

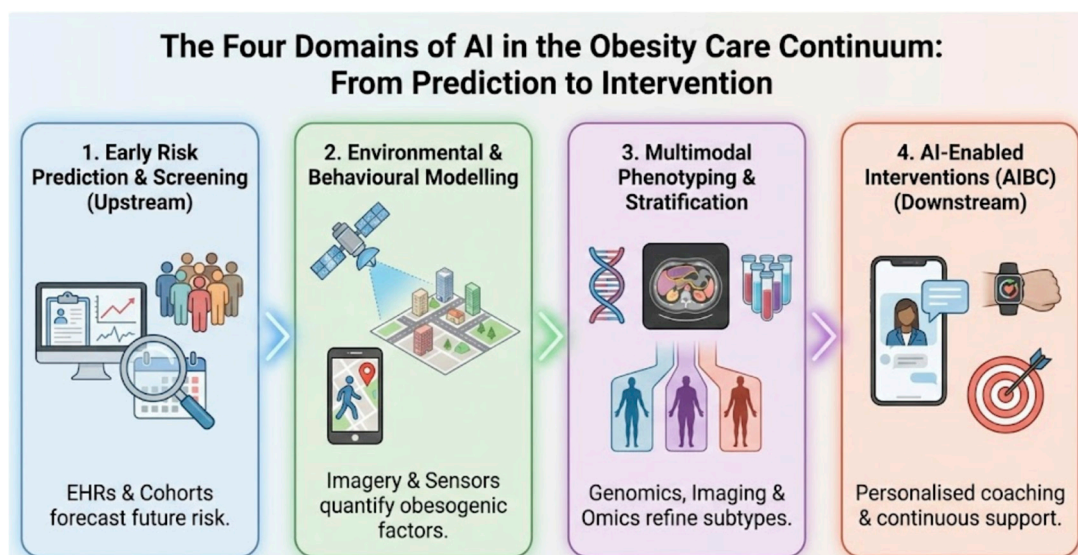
Titles and abstracts were screened for relevance, followed by full-text assessment of key articles. Evidence was synthesized thematically rather than quantitatively, with studies grouped into four domains: (1) early risk prediction and screening (EHR-based models and cohort analyses), (2) environmental and behavioural determinants (including image-based built-environment models), (3) multidimensional phenotyping and precision risk stratification (genomic and multimodal models), and (4) AI-enabled behavioural coaching and digital therapeutics, including implications for clinical governance. Priority was given to high-quality evidence (systematic reviews, large cohorts, and randomized trials) and to clinically actionable models with clear endpoints and validation approaches.

Given the narrative review design, no formal risk-of-bias tool was applied across all included studies. To mitigate selection bias and enhance interpretability, we preferentially included studies with transparent reporting of datasets, validation strategies (external validation where available), and clinically meaningful outcomes. Limitations include the potential for incomplete retrieval of relevant literature and heterogeneity in study designs, which precluded meta-analysis.

The authors acknowledge the use of an image generation model (Gemini 2.5 Flash Image / Nano Banana) to facilitate the conceptualization and initial drafting of the schematic illustrations presented in this review. All AI-assisted visual elements were subsequently critically reviewed, refined, and validated by the human authors to ensure scientific accuracy and alignment with the manuscript's content

### 3. Evidence Synthesis AI Applications Across the Obesity Care Continuum

The following synthesis organises current evidence into four clinically grounded domains. For each domain, we assess translational readiness with reference to validation design, endpoint selection, workflow fit, and governance requirements. The structure moves from upstream prediction to downstream delivery, reflecting the sequence in which AI tools encounter patients across an obesity care pathway, as illustrated in Figure 2.



**Figure 2.** The four domains of AI applications across the obesity care continuum: from prediction to intervention. Domain 1 (Early Risk Prediction and Screening, upstream): EHR-based and cohort-derived models forecast incident obesity and cardiometabolic complications. Domain 2 (Environmental and Behavioural Modelling): satellite and street-level imagery analysed by CNNs quantifies obesogenic environmental exposures at neighbourhood scale. Domain 3 (Multimodal Phenotyping and Stratification): integration of genomics, imaging, and multi-omics data delineates biologically meaningful obesity subtypes beyond BMI. Domain 4 (AI-Enabled Interventions, AIBC, downstream): personalised coaching platforms deliver continuous, adaptive behavioural support. Abbreviations: EHR, electronic health record; CNN, convolutional neural network; BMI, body mass index; AIBC, AI-enabled behavioural coaching.

### 3.1. Domain 1: Early Risk Prediction and Screening (EHR, Cohorts, Wearables as Inputs)

A central translational question is whether AI models can identify high risk trajectories early enough to enable preventive action before complications emerge. A recent narrative review in *Diagnostics* (2025) provides an upstream anchor by synthesizing three pillars of ML and deep learning applications in obesity: EHR based prediction of incident obesity, image based quantification of built environment correlates using satellite and street level data, and multimodal models integrating genetic and phenotypic variables for refined risk stratification and phenotyping. In the taxonomy adopted here, these pillars primarily populate the screening and stratification layers of the pathway and can subsequently be operationalized through clinical decision support and AI enabled behavioural coaching platforms that tailor intervention intensity and modality over time [4,5].

In relation to early screening, multiple studies have used electronic health records to predict incident obesity in children, including cohorts with up to approximately 11 million pediatric encounters. Algorithms such as random forests, gradient boosting and deep neural networks are trained on longitudinal anthropometric trajectories, perinatal factors, comorbid diagnoses, medication exposures and sociodemographic variables to identify children at high risk of developing obesity within prespecified prediction windows (for example, 1 to 5 years)[4]. These models are typically evaluated using hold out test sets or cross validation and achieve clinically useful discrimination across diverse health system datasets, supporting their potential use as screening tools to trigger targeted early life preventive interventions and more intensive follow up for those most likely to exceed BMI thresholds in later childhood[7].

Evidence from adult cohorts further supports the feasibility of AI based risk prediction. Within the Korean KNHANES programme, encompassing data from 21,100 adults, researchers tested multiple models to predict obesity risk[8]. Tree based methods such as random forest and XGBoost analyzed metabolic variables including triglycerides, liver enzymes and uric acid, with predictive performance varying by sex and age and highest accuracy observed in younger adults [9]. Reported

feature importance profiles suggested differences across subgroups, for example ALT ranking highly in men and triglycerides in women, indicating the potential value of stratified model development for personalised risk assessment. [9]

Interpretability tools are a practical requirement for clinical adoption, not merely a technical refinement. In NHANES-derived models combining metabolic and behavioural variables, SHAP analysis has identified daily step count, sedentary time, and fasting glucose as the strongest predictors of obesity risk—findings that translate directly into target-setting for lifestyle counselling. The ability to generate patient-level explanations alongside risk estimates substantially reduces the barrier to integrating ML outputs into consultation workflows.[11].

One advantage of predictive models is their ability to capture nonlinear interactions that may be overlooked in standard BMI-based assessments — an important consideration given that obesity pathophysiology extends well beyond simple energy balance [6]. For example, rising glucose in combination with low physical activity and prolonged sedentary exposure in early childhood may increase the probability of subsequent obesity over multi year horizons, suggesting a potential role for AI enabled screening integrated into electronic health records to identify children requiring intensified preventive measures, including behavioural and family based support.

Large-scale evidence synthesis further supports the temporal predictive capacity of ML models in obesity. Kalthori et al. conducted a systematic review including 10 studies across 411,000 participants with follow-up of one to five years, selected from an initial pool exceeding 6,000 publications. Inputs included age, sex, lifestyle variables, blood biomarkers and anthropometric measures, enabling machine learning algorithms to model nonlinear relationships and interacting factors that can be difficult to capture with traditional statistical approaches[12].

**Strength of evidence and remaining gaps.** The evidence base for EHR-based obesity risk prediction is the most mature among the four domains, supported by multiple large-cohort retrospective studies (n up to 11 million encounters), systematic review-level synthesis, and consistent demonstration of clinically useful discrimination (AUC 0.80–0.93 across age subgroups). Interpretability methods such as SHAP have been applied to generate patient-level explanations, strengthening the case for clinical adoption. However, critical gaps persist: almost all models rely on internal validation (hold-out or cross-validation) within the originating health system, and external validation across geographically or ethnically distinct populations remains rare. No prospective implementation study has yet demonstrated that embedding these models into clinical workflows reduces obesity incidence or improves downstream health outcomes. Additionally, the predominance of paediatric cohorts from the United States limits generalisability to adult populations and non-US health systems. Future work should prioritise prospective, multi-site validation with pre-registered clinical endpoints and equity-stratified performance reporting.

### 3.2. Domain 2: Environmental and Behavioural Determinants (Built Environment, Imagery)

The built environment is an established determinant of obesity risk, with systematic evidence demonstrating that environmental components — including access to recreational facilities, food environments, and neighbourhood walkability — are integral to effective prevention programmes [13]. Convolutional neural networks (CNNs) have been applied to satellite and street-level imagery to quantify such features at neighbourhood or census-tract level at scale [4]. In these studies, CNNs are trained on high-resolution images from platforms such as Google Street View or satellite providers to extract latent visual features related to land-use mix, building density, road structure, availability of green space, sidewalks and recreational facilities. These image-derived representations are then linked to population-level BMI or obesity prevalence data from surveys or health records, enabling identification of environmental configurations that correlate with higher or lower obesity burden. This approach supports scalable, standardised assessment of obesogenic and protective environments across large geographic areas, complementing labor-intensive environmental audits and informing urban planning and public health policy.

Recent advances have substantially scaled, refined, and integrated these CNN-based methods with advanced geospatial analytics, demonstrating their maturity as actionable public health tools [4]. A landmark 2025 JAMA Network Open study across 10 major US metropolitan areas (covering >1000 census tracts) employed deep CNN architectures to extract built environment features from both high-resolution satellite imagery and Google Street View panoramas, achieving 20–30% explained variance in neighborhood-level adult obesity prevalence—significantly outperforming traditional socioeconomic status (SES) indicators such as income, education, and deprivation indices. The models identified specific obesogenic configurations (e.g., high building density with low walkability) and enabled geospatial hotspot mapping via clustering algorithms, offering urban planners prioritized intervention zones with estimated impact on obesity reduction [15].

Similarly, a 2024 Missouri statewide geospatial analysis (published early 2025) integrated ResNet-50 CNN-extracted visual features from 10m-resolution satellite imagery with spatial econometric models (spatial lag/error specifications) and Local Indicators of Spatial Association (LISA), predicting obesity prevalence with high accuracy across urban, suburban, and rural census tracts. Key modifiable drivers included low vegetation density, single-lane road dominance, and fragmented green space, with rural areas showing distinct patterns of geographic isolation amplifying risk. The spatial models outperformed non-spatial baselines ( $R^2$  improvement of 15–25%), highlighting geographic autocorrelation as a critical confounder absent in traditional audits [14]

A 2025 multi-city Scientific Reports study spanning 19 diverse US urban areas (varying climate zones and demographics) further validated the approach using multi-modal CNN ensembles on Street View and Landsat satellite data, explaining ~25% of variance in both obesity prevalence and self-reported physical activity. Walkable streetscapes with  $\geq 30\%$  green coverage emerged as consistently protective ( $r = -0.35$  to  $-0.45$ ), while car-centric designs (wide roads, low mixed-use) correlated with higher BMI trajectories. The study demonstrated cross-city generalizability through transfer learning, reducing retraining needs for policy scalability [16]

Across these studies, CNN-based imagery analysis has demonstrated a consistent ability to identify specific, modifiable environmental configurations—low walkability, sparse green space, car-centric road design—that predict neighbourhood-level obesity prevalence at scale and with accuracy that manual audits cannot match. The practical implication is that area-level environmental risk scores derived from satellite and street-view data could serve as inputs for targeted public health resource allocation, informing where AIBC programmes or urban redesign investments are likely to have the greatest preventive impact. Rural-urban differences in risk patterns, as demonstrated in the Missouri analysis, suggest that these tools should be deployed with geographically stratified models rather than universal algorithms [16]

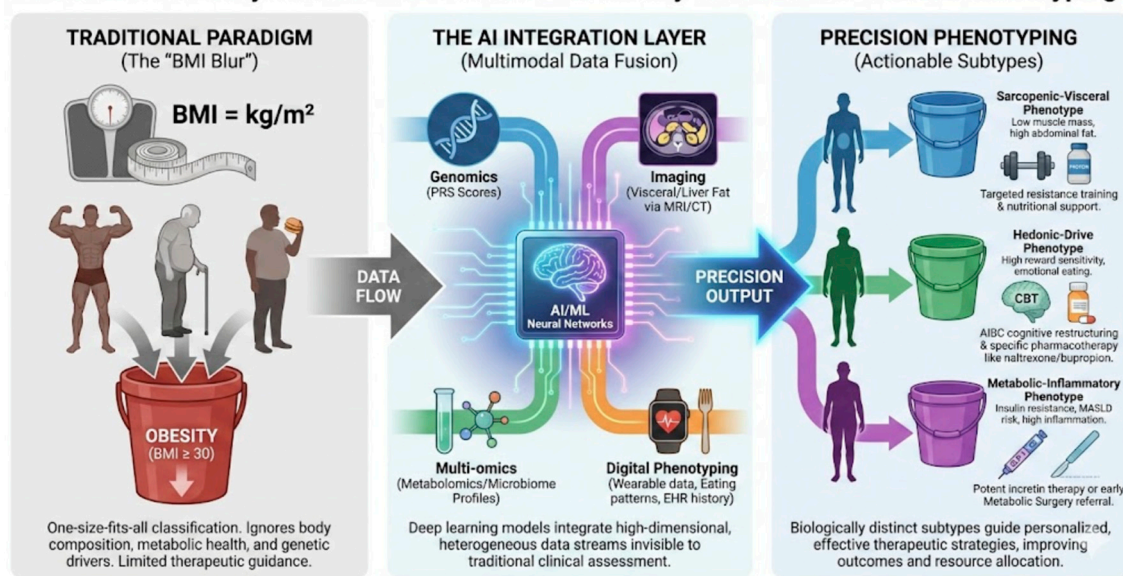
**Strength of evidence and remaining gaps.** Domain 2 evidence is anchored by several large-scale geospatial studies employing robust CNN architectures across diverse US metropolitan areas, achieving 20–30% explained variance in neighbourhood-level obesity prevalence and outperforming conventional SES indicators. The use of spatial econometric models and transfer learning across cities further strengthens methodological rigour. Nevertheless, these models operate exclusively at the population level; no study has yet linked area-level environmental risk scores to individual-level clinical decision-making or intervention outcomes. Geographic coverage remains confined to the United States, leaving it unclear whether identified obesogenic configurations (e.g., low walkability, sparse green space) replicate in European, Asian, or low- and middle-income settings where urbanisation patterns differ. Prospective studies evaluating whether CNN-derived environmental vulnerability indices can inform resource allocation decisions—and whether acting on such information reduces obesity prevalence—are needed.

### 3.3. Domain 3: Multimodal Phenotyping and Precision Risk Stratification (Genomics, Imaging, Omics)

Multimodal ML and deep learning models integrate genetic information, including single nucleotide polymorphisms and polygenic risk scores, with clinical, behavioural and environmental

covariates to refine individual risk estimation and delineate clinically meaningful obesity phenotypes. Methods such as regularized regression, tree based ensembles and neural networks combine hundreds to thousands of genetic variants with phenotypic data, such as BMI trajectories, metabolic markers and lifestyle factors [4]. (Figure 3)

### Evolution of Obesity Stratification: From BMI-Centricity to AI-Driven Precision Phenotyping



**Figure 3.** Evolution of obesity stratification: from BMI-centricity to AI-driven precision phenotyping. The left panel illustrates traditional BMI-based classification, which applies a single anthropometric threshold irrespective of body composition, metabolic health, or genetic background. The central panel depicts the AI integration layer, where deep learning models fuse multi-omics profiles (genomics, metabolomics, microbiome), imaging-derived measures of visceral adiposity (MRI/CT), and digital phenotyping data from wearables and EHR to generate high-dimensional feature representations. The right panel shows the resulting precision phenotypes—including sarcopenic-visceral, hedonic-drive, and metabolic-inflammatory subtypes—each associated with distinct pathophysiology and proposed targeted therapeutic strategies ranging from resistance training and nutritional support to GLP-1-based pharmacotherapy and cognitive-behavioural interventions. Abbreviations: BMI, body mass index; MRI, magnetic resonance imaging; CT, computed tomography; GLP-1, glucagon-like peptide-1; CBT, cognitive behavioural therapy; MASLD, metabolic dysfunction-associated steatotic liver disease; AIBC, AI-enabled behavioural coaching.

Across the studies reviewed, joint models generally outperform those using either genomic or clinical data alone for predicting obesity status or related cardiometabolic outcomes. Unsupervised clustering on integrated feature sets has also been used to identify subgroups with differing pathophysiological profiles and potentially differential response to interventions, providing a methodological basis for more personalised prevention and management strategies [4,17,18]. This advantage has been quantified systematically by Kim et al. (2025), who reported that multimodal AI models outperformed their unimodal counterparts in 91% of evaluated studies, with performance gains ranging from 6 to 33% depending on the combination of data modalities and clinical endpoint [19]. Complementary evidence from Hosseini et al. (2025) demonstrates that EHR data alone—including comorbidity trajectories, medication sequences, and longitudinal anthropometric patterns—can support deep phenotyping of obesity subtypes without requiring omics or imaging inputs, suggesting a lower-cost pathway to clinically actionable stratification in settings where multimodal data are unavailable [20].

Recent multimodal modelling efforts have further validated and extended this paradigm by demonstrating superior predictive performance and emerging clinical utility for treatment response stratification. A 2026 study in *npj Digital Medicine* developed an ML-derived polygenic risk score

(PRS) from UK Biobank data (n=482,700 individuals), integrating hundreds of thousands of SNPs with longitudinal phenotypic, lifestyle, and environmental covariates to forecast BMI trajectories across the life course; the joint model outperformed standard PRS by 12–18% in AUC for obesity incidence and persistence, capturing non-linear gene–environment interactions missed by univariate approaches. Similarly, a 2025 scoping review in *Advances in Nutrition* synthesized over 100 studies using regularized regression, tree-based ensembles, and neural networks to combine genetic variants with metabolomic, microbiome, and imaging data (e.g., visceral fat distribution from MRI/CT), consistently showing multimodal models superior to genomic or clinical silos (mean AUC gain 0.08–0.15) for cardiometabolic outcomes.[17,21]

Unsupervised clustering on these integrated datasets has yielded actionable phenotypes with differential intervention responses. Preliminary conference data from DDW 2025 reported ML phenotyping identifying novel obesity subtypes with distinct GLP-1 response profiles (e.g., metabolically healthy vs insulin-resistant visceral adiposity), validated in large cohorts for personalised pharmacotherapy matching (unpublished conference data). Likewise, Jia et al. (2025) in *Cell Reports Medicine* applied DDRTree unsupervised clustering to routine clinical data from 18,733 participants, identifying five metabolic phenotypes with distinct cardiovascular and type 2 diabetes risk profiles and proposed phenotype-specific treatment strategies, including GLP-1 prioritisation for subtypes characterised by high visceral adiposity [18]. These advances provide a methodological foundation for precision obesity care, transforming abstract risk scores into pathway-specific decision aids that inform whether patients require lifestyle intensification, early pharmacotherapy, or MBS referral (preliminary conference data)[18].

**Strength of evidence and remaining gaps.** Multimodal phenotyping models consistently demonstrate superior predictive performance over unimodal approaches—quantified at 6–33% improvement in a recent comprehensive review [19]—and unsupervised clustering has identified clinically meaningful subtypes with differential cardiometabolic risk profiles. The integration of ML-derived polygenic risk scores with lifestyle and environmental data in UK Biobank (n = 482,700) represents the largest-scale validation to date. However, this domain has the widest gap between methodological promise and clinical translation. No prospective trial has tested whether phenotype-guided treatment allocation improves outcomes compared with standard care. Most models are trained on data from European-ancestry biobanks, raising equity concerns for underrepresented populations. The high cost and limited availability of multi-omics and advanced imaging data restrict scalability, although EHR-based deep phenotyping approaches offer a lower-cost alternative. Validation against clinically actionable endpoints—such as differential response to GLP-1 RA versus cognitive-behavioural therapy by phenotype—should be a priority for the next generation of studies. A recent comprehensive review spanning mechanistic, clinical, and population-level AI applications in obesity further underscores the need for federated learning and privacy-preserving approaches to enable multi-site model development without centralising sensitive patient data [22]

#### 3.4. Domain 4: AI Enabled Lifestyle Interventions and Behavioural Coaching (AIBC)

Building on upstream prediction and stratification, Hallock et al. (2026) extend the role of AI in obesity medicine from identifying who is at risk to dynamically tailoring how care is delivered over time[23]. Whereas the Diagnostics review focuses on models trained on EHR, environmental imagery and genomic data to quantify and explain obesity risk, Hallock and colleagues describe how similar multimodal data streams can be integrated into AI enabled behavioural coaching platforms and clinical decision support systems that personalize intervention intensity, content and mode of delivery[4]. In this framework, outputs from predictive models, such as pediatric risk scores, environmental vulnerability indices or genetically informed phenotypes, serve as inputs for tools that adjust lifestyle coaching, inform choices across lifestyle modification, pharmacotherapy and metabolic or bariatric surgery, and support continuous monitoring and adaptation of treatment.

AI enabled behavioural coaching refers to digital interventions that use AI to deliver continuous, personalised lifestyle support for individuals with obesity. Systems may combine conversational

agents, just in time adaptive interventions, digital twin components and engagement mechanisms such as gamification[24]. Personalization is typically driven by multimodal data streams, including step counts, heart rate and sleep metrics from wearables, alongside self reported diet, mood and contextual factors. On this basis, platforms tailor the timing and frequency of prompts, the content and tone of messages and the specific tasks proposed to the user, aiming to sustain engagement and adherence over time [25,26]

Emerging clinical evidence suggests that AIBC interventions can achieve non inferior outcomes compared with standard care and may be associated with clinically meaningful weight loss, often defined as at least 5 percent of baseline body weight, alongside improvements in metabolic parameters including glycemic control, lipid profiles and blood pressure. However, the evidence base remains constrained by small samples, short follow up and heterogeneity in intervention design, target populations and outcome definitions, limiting generalizability and precluding firm conclusions about long term effectiveness or cost effectiveness [27]. A 2025 systematic review and meta-analysis of smartphone-based nutritional interventions confirmed modest but statistically significant effects on weight loss and body composition in adults with overweight or obesity, while underscoring that app-only interventions without additional AI personalisation or human support yield smaller effect sizes than hybrid approaches[28]. In addition, many AIBC solutions have been developed independently of formal clinical practice guidelines or clear regulatory frameworks for digital therapeutics, raising questions about guideline concordance, data governance and accountability.

Randomized evidence has begun to address these translational requirements [29]. In a trial embedded in the Diabetes Prevention Programme framework, an AI powered lifestyle intervention achieved non inferior weight loss and glycemic outcomes compared with standard human delivered coaching among adults at high risk of type 2 diabetes, many of whom had overweight or obesity[29]. The AI arm delivered automated, algorithmically tailored content and feedback, while the control arm followed the conventional DPP model with trained human coaches. Comparable outcomes support the feasibility of scaling behavioural programmes through AI while maintaining clinically meaningful effects [23].

**Strength of evidence and remaining gaps.** Domain 4 has produced the strongest direct evidence of patient benefit, including a landmark RCT demonstrating non-inferiority of AI-delivered coaching to human coaches within the Diabetes Prevention Programme framework (Mathioudakis et al., 2025). Additional mixed-methods and feasibility studies support the acceptability and short-term effectiveness of AI-assisted dietary and physical activity interventions. However, the evidence base remains thin in several respects: sample sizes are modest, follow-up rarely exceeds 12 months, and outcome definitions vary across studies, precluding quantitative synthesis. Cost-effectiveness data are absent. Critically, most AIBC platforms have been evaluated in high-income, digitally literate populations, leaving generalisability to low-resource settings undemonstrated. The extension of AI-driven coaching to clinically sensitive populations such as pregnant women remains at the protocol stage. Regulatory pathways for digital therapeutics in obesity are emerging (e.g., FDA clearance of Signos) but are not yet standardised across jurisdictions. Long-term, multi-site RCTs with equity-stratified endpoints and health-economic evaluation are the most pressing evidence needs.

### *3.5. Cross Cutting Synthesis: Translational Readiness, Guidelines, Regulatory Milestones and Deployment Signals*

The convergence of guideline endorsement, regulatory clearance, and commercial scaling that occurred between 2024 and 2026 represents a qualitative inflection point for AI in obesity medicine— one that moves the field beyond proof-of-concept and into questions of implementation fidelity, equity of access, and post-deployment governance. In late 2025, the World Health Organization issued its first global guideline on GLP-1 receptor agonists for obesity treatment, providing conditional recommendations for long-term pharmacotherapy in adults alongside intensive behavioural interventions as essential adjuncts. This landmark document explicitly recognizes the

need for comprehensive, multidisciplinary care integrating diet, physical activity, and professional support, directly validating the role of AI-enabled behavioural coaching (AIBC) as a scalable delivery mechanism for the behavioural pillar within personalised pathways combining lifestyle modification and pharmacotherapy[30].

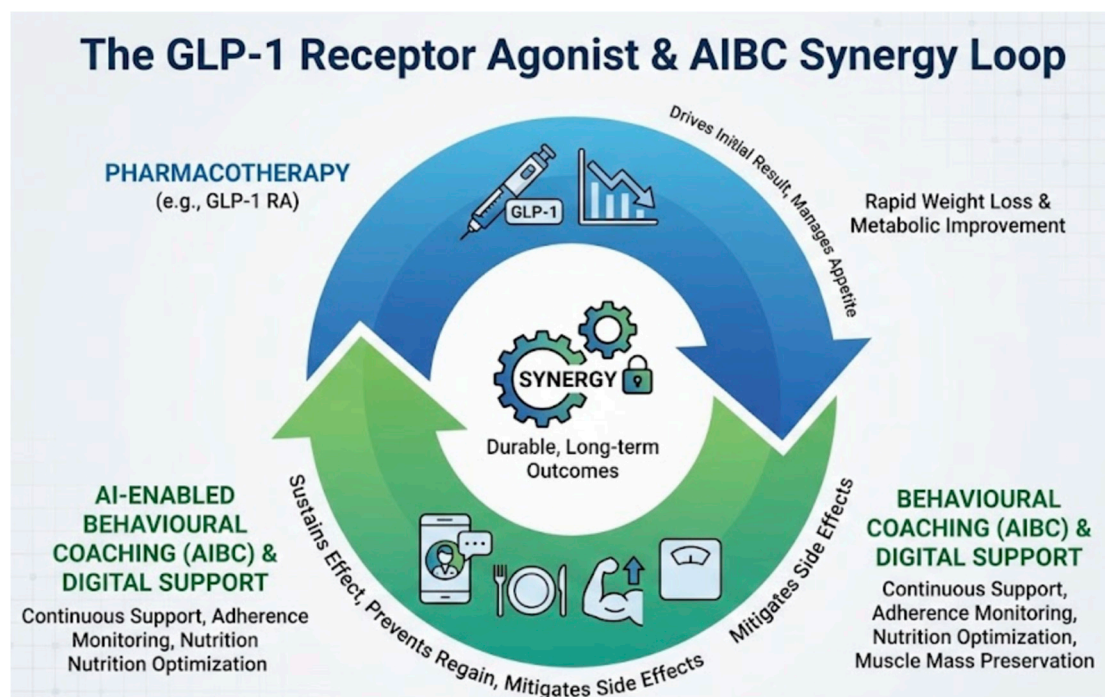
Similarly, the 2025 Obesity Algorithm update from the Obesity Medicine Association incorporates AI-driven prediction and monitoring tools into its decision trees for risk assessment, phenotyping, and therapy selection across lifestyle, pharmacotherapy, and metabolic/bariatric surgery (MBS) pathways[31]. Concurrently, the Gastroenterology consensus algorithm for obesity evaluation leverages the 2025 Lancet diagnostic criteria (clinical obesity plus individual risk factors), creating visual pathways where AI phenotyping and risk stratification tools fit seamlessly as enablers of escalated, personalised care[23,32]

Regulatory milestones further underscore clinical viability. In 2025, the FDA cleared Signos as the world's first AI-driven app paired with continuous glucose monitoring (CGM) for weight management, employing biomarker-informed nudges to optimize food and lifestyle choices based on real-time glucose responses [33]. These approvals establish a viable FDA pathway for AI obesity DTx, demonstrating regulatory feasibility while highlighting equity challenges (access/cost barriers) central to our ethical discussion.

At the commercial level, the integration of GLP-1 pharmacotherapy with continuous AI-enabled monitoring is moving from concept to product. Allurion Technologies' AllurionMeds platform—which pairs GLP-1 prescribing with AI-driven coaching and body composition tracking—exemplifies this model, targeting the clinically important problem of weight regain following GLP-1 discontinuation. Whether such integrated programmemes produce durable outcomes beyond the pharmacotherapy phase remains an open question requiring prospective data.(Figure 4)[34]. IQVIA's analysis notes the rise of FDA-cleared over-the-counter CGMs (Abbott Lingo, Dexcom Stelo) integrated with AI apps for glucose-guided obesity management, addressing adherence gaps in GLP-1 therapy. Reports like Obesity Management: Advances in 2025 confirm AI + GLP-1s as dominant 2025 innovations, with market growth in CGM/AI platforms[35]

The pace of this transition from research to practice is itself a governance risk: when deployment outstrips validation, inequities in access and algorithmic harm can become entrenched before surveillance mechanisms are in place. The evidence reviewed here suggests that clinical maturity is uneven across domains—strongest in EHR-based prediction and RCT-tested AIBC, thinner in multimodal phenotyping and pathway orchestration—and that translation decisions should be calibrated accordingly.

A comprehensive summary of the key studies reviewed across all four domains—including study design, population, sample size, data modality, AI method, primary outcome, explainability approach, validation status, and translational readiness level—is provided in Supplementary Table S1.



**Figure 4.** The GLP-1 receptor agonist and AI-enabled behavioural coaching (AIBC) synergy loop. GLP-1 pharmacotherapy drives rapid initial weight loss and metabolic improvement; AIBC provides continuous behavioural support, adherence monitoring, and muscle mass preservation, thereby mitigating common adverse effects and addressing the risk of weight regain following GLP-1 discontinuation. The synergistic interaction between pharmacological and digital therapeutic components is proposed as the dominant model for personalised obesity management in 2025–2026. Abbreviations: GLP-1 RA, glucagon-like peptide-1 receptor agonist; AIBC, AI-enabled behavioural coaching.

#### 4. Discussion

This narrative review demonstrates that AI applications in obesity care have moved beyond proof-of-concept, yet the four domains examined differ markedly in translational maturity—from implementation-ready EHR-based screening (Domain 1) and RCT-validated behavioural coaching (Domain 4), through scalable but population-level environmental modelling (Domain 2), to scientifically compelling but clinically unvalidated multimodal phenotyping (Domain 3). The domain-specific evidence profiles, including their respective strengths and remaining gaps, have been detailed in the preceding synthesis (Sections 3.1–3.4 and Supplementary Table S1). The discussion that follows therefore focuses on contextualising these findings against prior reviews, extracting clinical implications per domain, and addressing governance requirements for safe translation.

These findings extend and complement several recent reviews. Azmi et al. (2025) provided a comprehensive catalogue of AI techniques applied to obesity but focused primarily on methodological description rather than translational readiness or governance requirements [4]. Huang et al. (2025) offered a broader perspective on AI in obesity risk prediction and management but did not systematically evaluate the AIBC landscape, regulatory milestones, or implementation barriers [5]. The present review adds three contributions that were not addressed by these earlier works: first, a functional taxonomy that explicitly maps AI tools to their position and role within clinical obesity pathways, from upstream screening to downstream behavioural delivery; second, a domain-by-domain assessment of translational readiness, including validation design and endpoint maturity; and third, the integration of rapidly emerging evidence on AI-GLP-1 pharmacotherapy synergies, FDA regulatory clearances, and WHO guideline endorsements that post-date the search windows of prior reviews.

Importantly, the maturity gap between domains has practical implications for implementation sequencing. EHR-based screening tools and validated AIBC platforms are closest to deployment readiness and may yield the most immediate clinical returns, whereas multimodal phenotyping tools—though scientifically compelling—will require prospective validation in phenotype-guided treatment allocation trials before they can be recommended for routine use. This staged approach to adoption would allow health systems to accrue implementation experience and build governance infrastructure incrementally, rather than attempting simultaneous deployment across all AI modalities.

From a clinical standpoint, several domain-specific implications deserve emphasis. In Domain 1, the consistent ability of EHR-based models to identify children and adults on high-risk obesity trajectories 1–5 years in advance suggests a concrete use case: embedding validated risk scores into routine well-child visits and annual health checks to trigger referral to preventive programmes, analogous to cardiovascular risk calculators already in clinical use. The sex- and age-specific feature importance profiles observed in the KNHANES and NHANES analyses further suggest that a single universal model may be suboptimal; instead, stratified models calibrated to local demographic profiles are likely to yield better positive predictive values and reduce alert fatigue in primary care settings.

In Domain 2, the demonstration that CNN-derived environmental risk scores outperform conventional socioeconomic indicators for predicting neighbourhood-level obesity prevalence creates an opportunity for collaboration between public health agencies and urban planning authorities. Area-level environmental vulnerability indices could inform where to invest in walkability infrastructure, green space expansion, and AIBC programme deployment—linking AI-derived population intelligence to actionable policy levers.

In Domain 3, the identification of metabolically distinct obesity subtypes—such as the five phenotypes reported by Jia et al. (2025) with differential cardiovascular and type 2 diabetes risk—raises the possibility of phenotype-guided treatment selection. If prospective trials confirm that, for example, patients with high visceral adiposity and insulin resistance benefit preferentially from early GLP-1 RA initiation while those with a hedonic-drive phenotype respond better to cognitive-behavioural approaches combined with naltrexone/bupropion, the clinical impact would be substantial. However, until such trials are completed, phenotype-guided pathways should be considered hypothesis-generating rather than ready for clinical adoption. Notably, Xiao et al. (2026) demonstrated that the prognostic impact of a high ML-derived polygenic risk score was substantially attenuated by adherence to a healthy lifestyle, suggesting that AI-based genetic risk stratification should be framed as motivational rather than deterministic—reinforcing the link between upstream phenotyping (Domain 3) and downstream behavioural coaching (Domain 4)[21].

In Domain 4, the non-inferiority of AI-delivered coaching to human coaches in the DPP trial carries direct implications for scalability and reimbursement. If replicated in larger, longer-duration trials with diverse populations, these findings would support the inclusion of AI-based behavioural interventions in national coverage frameworks, potentially expanding access to evidence-based obesity care in underserved areas where trained human coaches are scarce. The emerging model of combining GLP-1 pharmacotherapy with continuous AI-enabled behavioural monitoring—as exemplified by the Signos FDA clearance and the Allurion platform—represents a paradigm in which pharmacological and digital therapeutic modalities are co-prescribed rather than offered sequentially, though prospective evidence of additive benefit remains limited.

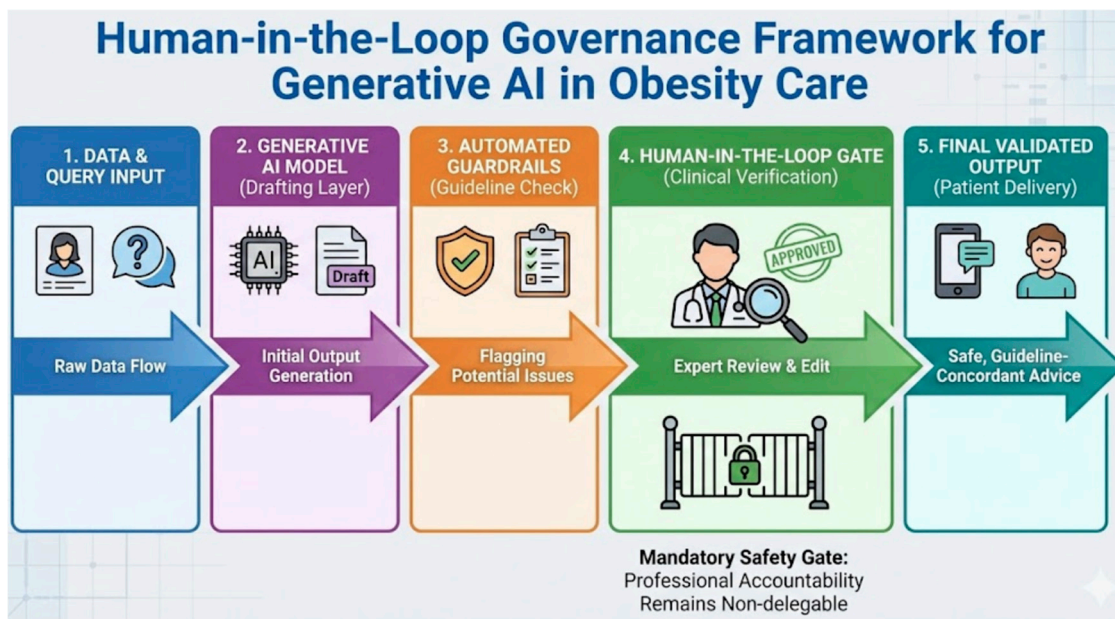
The clinical promise of AI in obesity care is inseparable from a set of governance challenges that deserve systematic attention. Chief among these are the risk of hallucination in generative systems deployed for patient education, the potential for algorithmic bias to amplify existing inequities in access to care, and the absence of clear accountability structures for patient-facing AI recommendations[5,10,23]. At the same time, many models are trained predominantly on data from high-income settings and majority populations, with limited representation of ethnic minorities and people living in low- and middle-income countries; this uneven data landscape increases the

likelihood that predictions of risk, treatment response or adherence will be less accurate—and potentially harmful—for precisely those groups already facing structural barriers to care [23,36–38]. In response to such concerns, the World Health Organization has articulated six core principles for ethical AI in health—respect for autonomy, promotion of human well-being and safety, transparency and explainability, responsibility and accountability, inclusiveness and equity, and sustainability—while professional bodies such as the AMA and AHA emphasize model explainability, continuous monitoring for harm and the preservation of clinician accountability[39]. However, current AIBC interventions in obesity rarely operationalize these frameworks in a systematic way, revealing a substantial gap between high-level guidance and day-to-day practice. Cultural and linguistic contexts must also be considered during implementation, as AI tools developed in high-income settings may embed assumptions that are poorly suited to the populations most affected by rising obesity prevalence. Translating these principles into practice requires attention to five specific requirements: prospective bias auditing across demographic subgroups before deployment; external validation in the local patient population, not only in the training cohort; defined human oversight protocols for all patient-facing outputs; explicit uncertainty communication to clinicians and patients; and post-deployment surveillance capable of detecting degradation in performance or new inequities as patient populations shift.

The extension of AI-enabled lifestyle interventions into clinically sensitive populations further amplifies the need for robust ethical and regulatory frameworks. A recent study of an AI-driven therapeutic lifestyle change (AI-TLC) programme in pregnant women with obesity illustrates this trajectory, using algorithmically tailored coaching to influence diet, physical activity and other health behaviors during gestation. While such interventions hold promise for improving maternal metabolic status and potentially reducing adverse pregnancy outcomes, they also operate in a context where the stakes are uniquely high: any unintended effects of algorithmic misclassification, inappropriate advice or unrecognized bias may impact not only the pregnant individual but also fetal development and long-term offspring health. When viewed alongside broader proposals for AI-enabled behavioural coaching and precision phenotyping in obesity, these findings suggest that regulatory expectations for validation, safety monitoring, transparency and human oversight should be at least as stringent for AI systems targeting high-risk groups—such as pregnant women—as for conventional pharmacological or surgical interventions, including clear standards for informed consent, data governance and equitable access[23,40].

Parity between AI coaching and human-delivered DPP has direct implications for reimbursement: digital therapeutics that demonstrate non-inferiority in RCTs should be evaluated under frameworks analogous to those applied to pharmacological or procedural interventions, including requirements for post-marketing surveillance and equity monitoring[29].

To mitigate the risk of generative hallucinations, the 'Human-in-the-loop' (HITL) architecture must be operationalized as a mandatory safety gate[41]. In this framework, LLM-generated patient education materials are not delivered directly; instead, they serve as 'drafting layers' for multidisciplinary teams. Professional accountability remains non-delegable; thus, AI tools should be viewed as 'clinical co-pilots' that synthesize data while the final therapeutic signature remains human-verified [42] as operationalised in the HITL framework illustrated in Figure 5



**Figure 5.** Human-in-the-loop (HITL) governance framework for generative AI in obesity care. The five-stage pipeline proceeds from (1) raw data and query input through (2) generative AI model output (drafting layer), (3) automated guardrails performing guideline concordance checks, (4) the mandatory human-in-the-loop safety gate at which clinicians review, edit, and approve AI-generated content, to (5) final validated patient-facing output. Professional accountability is non-delegable at Stage 4, ensuring that AI functions as a clinical co-pilot rather than an autonomous decision-maker. This architecture is proposed as a minimum governance standard for patient-facing AI applications in obesity management, consistent with WHO ethical AI principles and regulatory expectations for digital therapeutics. Abbreviations: HITL, human-in-the-loop; AI, artificial intelligence; LLM, large language model; WHO, World Health Organization.

This review is limited by the narrative design, which does not provide exhaustive retrieval or formal risk-of-bias assessment across all included studies. The heterogeneity of populations, data modalities, model architectures and outcome definitions limits comparability and precludes quantitative synthesis. In addition, publication bias and selective reporting of model performance may inflate apparent effectiveness, particularly in early-stage digital interventions. A further limitation is the rapid pace of the field: several studies cited here were published as preprints or conference abstracts and have not yet undergone full peer review, which should be considered when interpreting reported effect sizes. Additionally, the predominance of studies from high-income settings—particularly the United States, South Korea, and the United Kingdom—limits the generalisability of conclusions to low- and middle-income countries where obesity prevalence is rising most rapidly and health system infrastructure for AI deployment differs fundamentally. The review did not formally assess reporting quality of included AI studies against frameworks such as TRIPOD or CONSORT-AI, which would have strengthened the appraisal of model validation claims. Future systematic reviews incorporating formal quality assessment and meta-analysis will be essential to consolidate the evidence base as it matures.

## 5. Conclusions

Across the obesity care continuum, AI is increasingly capable of supporting earlier identification of high-risk trajectories, refining phenotypes beyond BMI, and scaling behavioural support through AI-enabled coaching and digital therapeutics. The most clinically mature signals are emerging in EHR-based risk prediction and selected digital lifestyle interventions, while multimodal phenotyping and pathway-level orchestration remain promising but unevenly validated. Translation into routine care will require external validation on local populations, explicit bias auditing, and human-in-the-

loop oversight, particularly for patient-facing generative systems vulnerable to hallucinations and guideline discordance. Future research should prioritize long-term effectiveness, cost-effectiveness, and implementation studies that measure equity and safety outcomes alongside weight and metabolic endpoints. When embedded within governance frameworks and multidisciplinary care models, AI has the potential to operationalize genuinely personalised obesity treatment pathways from screening to sustained management.

**Supplementary Materials:** The following supporting information can be downloaded at the website of this paper posted on Preprints.org.

**Funding:** This research received no external funding.

**Institutional Review Board Statement:** Not applicable. This study is a narrative review of previously published literature and did not involve human participants, animals, or the collection of personal data.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** No new data were created or analysed in this study. Data sharing is not applicable to this article.

**Acknowledgments:** The authors acknowledge the use of an image generation model (Gemini 2.5 Flash Image/Nano Banana) to facilitate the conceptualization and initial drafting of the schematic illustrations presented in this review. All AI-assisted visual elements were subsequently critically reviewed, refined, and validated by the human authors to ensure scientific accuracy and alignment with the manuscript's content.

**Conflicts of Interest:** The authors declare no conflicts of interest.

## Abbreviations

The following abbreviations are used in this manuscript:

AI, artificial intelligence;  
AIBC, AI-enabled behavioural coaching;  
AUC, area under the receiver operating characteristic curve;  
BMI, body mass index;  
CDSS, clinical decision support system;  
CGM, continuous glucose monitoring;  
CNN, convolutional neural network;  
CT, computed tomography;  
DPP, Diabetes Prevention Programme;  
DTx, digital therapeutics;  
DXA, dual-energy X-ray absorptiometry;  
EHR, electronic health record;  
FDA, US Food and Drug Administration;  
GLP-1 RA, glucagon-like peptide-1 receptor agonist;  
HITL, human-in-the-loop;  
JITAI, just-in-time adaptive intervention;  
KNHANES, Korea National Health and Nutrition Examination Survey;  
LMIC, low- and middle-income countries;  
LLM, large language model;  
MBS, metabolic and bariatric surgery;  
ML, machine learning;  
DL, deep learning;  
MRI, magnetic resonance imaging;  
NHANES, National Health and Nutrition Examination Survey;  
PRS, polygenic risk score;

RCT, randomised controlled trial;  
SES, socioeconomic status;  
SHAP, SHapley Additive exPlanations;  
SNP, single nucleotide polymorphism;  
WHO, World Health Organization;  
XGBoost, extreme gradient boosting.

## References

1. Wójcik S, Rulkiewicz A, Pruszczyk P, Lisik W, Poboży M, Domienik-Karłowicz J. Beyond ChatGPT: What does GPT-4 add to healthcare? The dawn of a new era. *Cardiol J.* 2023;30(6):1018-1025. doi: 10.5603/cj.97515. Epub 2023 Oct 13. PMID: 37830256; PMCID: PMC10713213.
2. Akbarialiabad H, Pasdar A, Murrell DF et al. Enhancing randomized clinical trials with digital twins. *npj Syst Biol Appl.* 2025;11:110. doi: 10.1038/s41540-025-00592-0.
3. Wójcik S, Rulkiewicz A, Pruszczyk P, Lisik W, Poboży M, Domienik-Karłowicz J. Reshaping medical education: Performance of ChatGPT on a PES medical examination. *Cardiol J.* 2024;31(3):442-450. doi: 10.5603/cj.97517. Epub 2023 Oct 13. PMID: 37830257; PMCID: PMC11229801.
4. Azmi S, Kunnathodi F, Alotaibi HF, Alhazzani W, Mustafa M, Ahmad I, Anvarbatcha R, Lytras MD, Arafat AA. Harnessing Artificial Intelligence in Obesity Research and Management: A Comprehensive Review. *Diagnostics (Basel).* 2025 Feb 6;15(3):396. doi: 10.3390/diagnostics15030396. PMID: 39941325; PMCID: PMC11816645.
5. Huang L, Huhulea EN, Abraham E, Bienenstock R, Aifuwa E, Hirani R, Schulhof A, Tiwari RK, Etienne M. The Role of Artificial Intelligence in Obesity Risk Prediction and Management: Approaches, Insights, and Recommendations. *Medicina (Kaunas).* 2025 Feb 19;61(2):358. doi: 10.3390/medicina61020358. PMID: 40005474; PMCID: PMC11857386.
6. Hall KD, Farooqi IS, Friedman JM, Klein S, Loos RJF, Mangelsdorf DJ, O'Rahilly S, Ravussin E, Redman LM, Ryan DH, Speakman JR, Tobias DK. The energy balance model of obesity: beyond calories in, calories out. *Am J Clin Nutr.* 2022 May 1;115(5):1243-1254. doi: 10.1093/ajcn/nqac031. PMID: 35134825; PMCID: PMC9071483.
7. Gupta M, Phan TT, Bunnell HT, Beheshti R. Obesity Prediction with EHR Data: A deep learning approach with interpretable elements. *ACM Trans Comput Healthc.* 2022 Jul;3(3):32. doi: 10.1145/3506719. Epub 2022 Apr 7. PMID: 35756858; PMCID: PMC9221869.
8. Jeong S, Choi YJ. Investigating the Influence of Heavy Metals and Environmental Factors on Metabolic Syndrome Risk Based on Nutrient Intake: Machine Learning Analysis of Data from the Eighth Korea National Health and Nutrition Examination Survey (KNHANES). *Nutrients.* 2024 Mar 2;16(5):724. doi: 10.3390/nu16050724. PMID: 38474852; PMCID: PMC10934821.
9. Jeon J, Lee S, Oh C. Age-specific risk factors for the prediction of obesity using a machine learning approach. *Front Public Health.* 2023 Jan 17;10:998782. doi: 10.3389/fpubh.2022.998782. PMID: 36733276; PMCID: PMC9887184.
10. Khera R, Butte AJ, Berkwits M, Hswen Y, Flanagan A, Park H, Curfman G, Bibbins-Domingo K. AI in Medicine—JAMA's Focus on Clinical Outcomes, Patient-Centered Care, Quality, and Equity. *JAMA.* 2023;330:818-820. doi: 10.1001/jama.2023.15481.
11. Netayawijit P, Chansanam W, Sorn-In K. Interpretable Machine Learning Framework for Diabetes Prediction: Integrating SMOTE Balancing with SHAP Explainability for Clinical Decision Support. *Healthcare (Basel).* 2025 Oct 14;13(20):2588. doi: 10.3390/healthcare13202588. PMID: 41154264; PMCID: PMC12563896.
12. Niakan Kalhori SR, Najafi F, Hasannejadasl H, Heydari S. Artificial intelligence-enabled obesity prediction: A systematic review of cohort data analysis. *Int J Med Inform.* 2025 Apr;196:105804. doi: 10.1016/j.ijmedinf.2025.105804. Epub 2025 Jan 24. PMID: 39870016.
13. Cauchi D, Glonti K, Petticrew M, Knai C. Environmental components of childhood obesity prevention interventions: An overview of systematic reviews. *Obes Rev.* 2016;17:1116-1130. doi: 10.1111/obr.12441.

14. Dahu BM, Khan S, Toubal IE, Alshehri M, Martinez-Villar CI, Ogundele OB, Sheets LR, Scott GJ. Geospatial Modelling of Deep Neural Visual Features for Predicting Obesity Prevalence in Missouri: Quantitative Study. *JMIR AI*. 2024 Dec 17;3:e64362. doi: 10.2196/64362. PMID: 39688897; PMCID: PMC11688583.
15. Chen Z, Zhang T, Dazard JE, Ponnana SR, Dong W, Moorthy S, Sirasapalli SK, Khraishah H, Deo S, Rajagopalan S, Al-Kindi S. AI-Enhanced Analysis of Built Environment Imagery and Neighborhood Obesity in US Cities. *JAMA Netw Open*. 2025 Sep 2;8(9):e2534612. doi: 10.1001/jamanetworkopen.2025.34612. Erratum in: *JAMA Netw Open*. 2025 Dec 1;8(12):e2554305. PMID: 41026491; PMCID: PMC12485642.
16. Ghorbany S, Hu M, Yao S et al. Data driven assessment of built environment impacts on urban health across United States cities. *Sci Rep*. 2025;15:19998. doi: 10.1038/s41598-025-04567-3.
17. Vahid F, Loyola-Leyva A, Tur J, Bouzas C, Devaux Y, Malisoux L, Garcia S, De Carvalho M, Ródenas-Munar M, Turner J, Lamy E, Perez-Jimenez M, Ravn-Haren G, Andersen R, Forberger S, Nagrani R, Onorati MG, Bonetti GG, Rodrigues D, Bohn T. Multimodal (Bio)Markers and Risk of Obesity – A Comprehensive Scoping Review. *Adv Nutr*. 2026;17(2):100579. doi: 10.1016/j.advnut.2025.100579.
18. Jia X, Lin H, Ding Y, Hu C, Wang S, Li M, Xu Y, Xu M, Huang F, Shen F, Gu X, Mu Y, Chen L, Zeng T, Shi L, Su Q, Yu X, Yan L, Qin G, Wan Q, Chen G, Tang X, Gao Z, Hu R, Luo Z, Qin Y, Chen L, Hou X, Huo Y, Li Q, Wang G, Zhang Y, Liu C, Wang Y, Wu S, Zhu Y, Yang T, Deng H, Zhao J, Zhang Y, Xu X, Wei H, Zheng J, Wang T, Zhao Z, Ning G, Chen Y, Wang W, Bi Y, Lu J. Phenotyping obesity through a two-dimensional tree structure reveals cardiometabolic heterogeneity. *Cell Rep Med*. 2025 Nov 18;6(11):102372. doi: 10.1016/j.xcrm.2025.102372. Epub 2025 Sep 25. PMID: 41005299; PMCID: PMC12711682.
19. Kim JH, Lee H, Hwang J, Yon DK. Multimodal and Multidimensional Artificial Intelligence Technology in Obesity. *J Obes Metab Syndr*. 2025;34(4):394-404. doi: 10.7570/jomes25035. PMID: 40922673.
20. Ruan X, Lu S, Wang L, Wen A, Murali S, Liu H. Deep Phenotyping of Obesity: Electronic Health Record-Based Temporal Modeling Study *J Med Internet Res* 2025;27:e70140 URL: <https://www.jmir.org/2025/1/e70140> DOI: 10.2196/70140
21. Xiao L, Liang S, Zeng L et al. A machine learning-derived polygenic risk score reveals that healthy lifestyle counteracts obesity-related mortality. *npj Digit Med*. 2026;9:72. doi: 10.1038/s41746-025-02314-6.
22. Wang MH. Artificial Intelligence Across the Obesity Continuum: From Mechanistic Insights to Global Precision Prevention and Therapy. *Obesity (Silver Spring)*. 2026 Feb;34(2):294-316. doi: 10.1002/oby.70095. Epub 2025 Dec 10. PMID: 41371937.
23. Hallock RK, Griebeler ML, Burguera B, Cabandugama P. Transforming obesity medicine with artificial intelligence: personalization, precision and the path ahead. *Obes Endocrinol*. 2026;2(1):wjaf018. doi: 10.1093/obendo/wjaf018.
24. Kotov A, Idalski Carcone A, Towner E. Neural Conversational Agent for Weight Loss Counseling: Protocol for an Implementation and Feasibility Study. *JMIR Res Protoc*. 2024 Sep 20;13:e60361. doi: 10.2196/60361. PMID: 39303273; PMCID: PMC11452760.
25. Nahum-Shani I, Smith SN, Spring BJ, Collins LM, Witkiewitz K, Tewari A, Murphy SA. Just-in-Time Adaptive Interventions (JITAs) in Mobile Health: Key Components and Design Principles for Ongoing Health Behavior Support. *Ann Behav Med*. 2018 May 18;52(6):446-462. doi: 10.1007/s12160-016-9830-8. PMID: 27663578; PMCID: PMC5364076.
26. Chew Hsj, Chew NW, Loong SSE, Lim SL, Tam WSW, Chin YH, Chao AM, Dimitriadis GK, Gao Y, So JBY, Shabbir A, Ngiam KY. Effectiveness of an Artificial Intelligence-Assisted App for Improving Eating Behaviors: Mixed Methods Evaluation. *J Med Internet Res*. 2024;26:e46036. doi: 10.2196/46036.
27. Noh E, Won J, Jo S, Hahm DH, Lee H. Conversational Agents for Body Weight Management: Systematic Review. *J Med Internet Res*. 2023;25:e42238. doi: 10.2196/42238.
28. Pujia C, Ferro Y, Mazza E, Maurotti S, Montalcini T, Pujia A. The Role of Mobile Apps in Obesity Management: Systematic Review and Meta-Analysis. *J Med Internet Res*. 2025;27:e66887. doi: 10.2196/66887. PMID: 40327853.
29. Mathioudakis N, Lalani B, Abusamaan MS, Alderfer M, Alver D, Dobs A, Kane B, McGready J, Riekert K, Ringham B, Shehadeh A, Vandi F, Wanigatunga AA, Zade D, Maruthur NM; AI-DPP Study Group. An AI-Powered Lifestyle Intervention vs Human Coaching in the Diabetes Prevention Programme: A

- Randomized Clinical Trial. *JAMA*. 2025 Dec 16;334(23):2079-2089. doi: 10.1001/jama.2025.19563. PMID: 41144242; PMCID: PMC12560030.
30. World Health Organization. WHO guideline on the use of GLP-1 receptor agonists for the management of obesity. Geneva: WHO; 2025. Available at: <https://www.who.int/news/item/01-12-2025-who-issues-global-guideline-on-the-use-of-glp-1-medicines-in-treating-obesity> [Accessed 05 Feb 2026].
  31. Obesity Medicine Association. Obesity Algorithm 2025. Denver: OMA; 2025. Available at: <https://obesitymedicine.org/resources/obesity-algorithm/> [Accessed 05 Feb 2026].
  32. Nadolsky K, Garvey WT, Agarwal M, Bonnacaze A, Burguera B, DeGeeter Chaplin M, Griebeler ML, Harris SR, Schellinger JN, Simonetti J, Srinath R, Yumuk V. American Association of Clinical Endocrinology Consensus Statement: Algorithm for the Evaluation and Treatment of Adults with Obesity/Adiposity-Based Chronic Disease – 2025 Update. *Endocr Pract*. 2025;31(11):1351-1394. doi: 10.1016/j.eprac.2025.07.017.
  33. US Food and Drug Administration. FDA clears Signos AI-driven glucose monitoring for weight management [news release]. 20 Aug 2025. Available at: <https://www.cnbc.com/2025/08/20/fda-approves-signos-glucose-monitoring-for-weight-loss.html> [Accessed 05 Feb 2026].
  34. Allurion Technologies. AllurionMeds: AI-driven coaching with GLP-1 prescribing [press release]. 2025.
  35. IQVIA Institute for Human Data Science. Obesity and Weight Management 2025 Report. Durham, NC: IQVIA; 2025.
  36. Otokiti AU, Shih HJ, Williams KS. Gender and racial bias unveiled: clinical artificial intelligence (AI) and machine learning (ML) algorithms are fanning the flames of inequity. *Oxf Open Digit Health*. 2025;3:oqaf027. doi: 10.1093/oodh/oqaf027.
  37. Joseph J. Algorithmic bias in public health AI: a silent threat to equity in low-resource settings. *Front Public Health*. 2025 Jul 23;13:1643180. doi: 10.3389/fpubh.2025.1643180. PMID: 40771228; PMCID: PMC12325396.
  38. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*. 2019 Oct 25;366(6464):447-453. doi: 10.1126/science.aax2342. PMID: 31649194.
  39. World Health Organization. Ethics and governance of artificial intelligence for health: WHO guidance. Geneva: WHO; 2021 (updated 2023). Available at: <https://www.who.int/publications/i/item/9789240029200> [Accessed 05 Feb 2026].
  40. Wang X, Bai Y, Li W, Yang C, Zhang L, Zhu H, Bao R, Jiang Y, Wang F, Wang H. Effect of artificial intelligence driven therapeutic lifestyle changes (AI-TLC) intervention on health behavior and health among obesity pregnant women in China: a randomized controlled trial protocol. *Front Public Health*. 2025 May 12;13:1580060. doi: 10.3389/fpubh.2025.1580060. PMID: 40421363; PMCID: PMC12104060.
  41. Crigger E, Reinbold K, Hanson C, Kao A, Blake K, Irons M. Trustworthy Augmented Intelligence in Health Care. *J Med Syst*. 2022 Jan 12;46(2):12. doi: 10.1007/s10916-021-01790-z. PMID: 35020064; PMCID: PMC8755670.
  42. Huang L, Yu W, Ma W, Zhong W, Feng Z, Wang H, Chen Q, Peng W, Feng X, Qin B, Liu T. A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *ACM Trans Inf Syst*. 2025;43(2):Article 42. doi: 10.1145/3703155.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.