children

MDPI

Article

Bridging Data, Semantics, and Clinical Reasoning: A Knowledge Graph Framework for Pediatric Obstructive Sleep Apnea

James D. Geyer [1,2,3,*], Jiaqi Gong [3], Paul G. Cox [4], Randi J. Henderson-Mitchell [1], Camilo R. Gomez [5], Adnan I. Qureshi [5], Shelby G. Branch [1], Sophia R. Geisser [6] and Paul R. Carney [7]

1 Institute for Rural Health Research, University of Alabama, Tuscaloosa, AL 35487, USA
2 University Medical Center Neurology and Sleep Medicine, University of Alabama, Tuscaloosa, AL 35487, USA
3 Alabama Center for the Advancement of AI, Department of Computer Science, University of Alabama, Tuscaloosa, AL 35487, USA
4 Evimero, LLC, Tuscaloosa, AL 35406, USA
5 Department of Neurology, University of Missouri School of Medicine, Columbia, MO 65212, USA
6 Department of Psychology, University of Alabama, Tuscaloosa, AL 35487, USA
7 Sunstone Health, Boston, MA 02101, USA
* Correspondence: jdgeyer@ua.edu

Highlights

What are the main findings?

- Current artificial intelligence for pediatric sleep medicine is limited by fragmented data, a lack of clinical explainability, and failure to account for age-specific developmental differences.
- A proposed knowledge graph framework successfully integrates multimodal data including polysomnography, electronic health records, and clinical narratives, into a unified reasoning system.

What are the implications of the main findings?

- This framework enables the development of interpretable clinical decision tools that provide clinicians with auditable evidence for personalized screening and post-surgical risk assessments.
- Standardizing pediatric sleep data facilitates large-scale research collaborations and enables the projection of specialized care into underserved and rural regions.

Abstract

Background/Objectives: Pediatric obstructive sleep apnea (OSA) is a complex disorder with a variable presentation and often challenging diagnostic testing. The history and physical examination in pediatric OSA frequently differ from those in adults. The treatment options are multifaceted and must be tailored to the individual patient. Artificial intelligence (AI) modalities currently employed in pediatric sleep medicine face several important limitations: modality fragmentation, lack of explainability, and limited semantic integration. Method: Our team proposes a new vision for AI and pediatric sleep medicine. This platform is based on a knowledge graph (KG) framework integrating structured and unstructured data to enable reasoning, personalization, and clinical decision support. Results: This framework represents a conceptual architecture; it has not yet been empirically implemented, and the use cases described herein are illustrative of its intended capabilities. Components of the infrastructure developed for similar applications have been successfully implemented. The quantitative feasibility pilot KG represented 100% multimodal data with >90% semantic completeness. Conclusions: Fully realized and deployed into the clinical

Check for updates

Academic Editor: Jun Kohyama

Received: 18 March 2026
Revised: 20 April 2026
Accepted: 22 April 2026
Published: 27 April 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

space, this pediatric OSA KG system will enhance tertiary care programs and help project tertiary-level pediatric care into underserved regions.

Keywords: pediatric obstructive sleep apnea; artificial intelligence; knowledge graph

1. Introduction

Obstructive sleep apnea (OSA) is defined by repeated episodes of partial or complete upper-airway obstruction during sleep. These episodes lead to disrupted sleep and abnormalities in gas exchange. The patient may or may not snore and is often unaware of the breathing changes [1]. The diagnosis of pediatric OSA is confirmed with overnight polysomnography (PSG), as recommended by the American Academy of Pediatrics (AAP), [2] American Academy of Sleep Medicine (AASM) [3], and European Respiratory Society (ERS) [4]; PSG measures the apnea-hypopnea index (AHI), defined as the number of apneas and hypopneas per hour of sleep. In children under 13 years of age, an AHI above 1 respiratory event/hour is considered abnormal. AHI classifies pediatric OSA severity into mild (1–4/h), moderate (5–9/h), and severe ($\geq$10/h). These thresholds are lower than those for adults and reflect children's greater sensitivity to airway obstruction. Adult OSA severity, which may be appropriate for assessment of adolescent OSA, is classified by AHI into mild (5–15/h), moderate (15–30/h), and severe ($\geq$30/h) [5]. The AAP advises that primary care providers screen for snoring and other OSA symptoms at routine health visits. This screening is compromised by the complexity of history and physical examination findings in sleep disorders. A thorough and optimized OSA evaluation is needed for any child with snoring or other symptoms before considering therapeutic interventions like surgery or continuous positive airway pressure (CPAP).

Epidemiology: OSA is common in childhood. Studies report that approximately 2–5% of the pediatric population has OSA. The peak prevalence occurs in early childhood, typically in ages 2–8 years. This is often secondary to enlarged tonsils and adenoids [6,7]. In adolescents, OSA remains an important issue, but the underlying etiologies shift, such as obesity [2]. The prevalence of OSA in older adolescents (16–19 years) is approximately 4%, with obesity and male sex as significant risk factors in this age group [8]. In obese teens, OSA can be strikingly common, potentially 40 to 60% [9]. Unlike in younger children, adolescent OSA more closely resembles adult patterns in epidemiology, symptomatology, physical findings, and polysomnographic findings. During childhood, OSA affects boys and girls roughly equally in early years. However, adolescent males with additional risk factors are more likely to have OSA [10]. Overall, pediatric OSA is underdiagnosed. Increased clinician awareness and screening are needed to identify affected children early.

Clinical Features and Symptoms: The nighttime symptoms of pediatric OSA often alert families to the problem. Loud snoring is common but not required and is frequently accompanied by observed pauses in breathing or gasping, labored breathing, restless sleep with frequent position changes, night sweats, or even bedwetting (nocturnal enuresis). Some children with OSA adopt unusual sleeping positions (e.g., neck hyperextended) to keep the airway open. Mouth breathing and dry mouth are also common due to chronic airway obstruction. By morning, children with OSA may wake with headaches or a dry throat. During the daytime, the classic adult symptom of excessive daytime sleepiness is less common in young children; instead, many pediatric patients with OSA exhibit behavioral and cognitive symptoms. Overt sleepiness during the day is rare in younger children with OSA and tends to be seen more in adolescents or in obese children. Rather than appearing lethargic, school-aged children with untreated OSA may struggle with behavioral and

neurocognitive issues that may closely resemble attention-deficit/hyperactivity disorder (ADHD) [11,12], increasing the risk of an incorrect diagnosis of ADHD and unnecessary stimulant therapy.

Against this backdrop, artificial intelligence (AI) has been explored to improve the detection, characterization, and management of pediatric OSA. In polysomnography (PSG) signal processing, convolutional and recurrent neural networks, and anomaly-detection methods have been developed to automate event detection, sleep staging, and arousal identification from multichannel physiological signals. These systems can reduce scoring time and inter-scorer variability, and they can unmask subtle temporal patterns (e.g., REM-predominant obstruction, positional dependence, or hypoventilation signatures) that inform risk and therapy selection. Beyond PSG, supervised machine-learning models using features extracted from the electronic health record (EHR) (e.g., demographics, anthropometrics, comorbid diagnoses, medications, and vital signs) have been trained to prioritize referrals, predict study positivity, or anticipate residual disease after surgery. Questionnaires such as the Pediatric Sleep Questionnaire (PSQ) [13] and fatigue/sleepiness scales have been mined with elastic-net, gradient-boosting, and calibration frameworks to refine screening thresholds and triage strategies. More recently, large language models (LLMs) have been utilized to summarize narrative sleep reports, harmonize impressions across providers, and answer clinician queries using textual context.

Despite these advances, current AI efforts are siloed by modality and task and therefore struggle to generalize across settings and populations. PSG-only models often ignore comorbid conditions and contextual factors (e.g., craniofacial morphology, oropharyngeal anatomy, allergic rhinitis, asthma) that shape risk and outcomes [14,15]. EHR-based risk scores can inherit coding biases, lack temporal granularity, and underperform in health systems with different documentation practices [16,17]. Questionnaire-centric models are sensitive to language, culture, and caregiver education [18]. Across modalities, explainability remains limited: many systems provide predictions without mechanistic pathways that clinicians can validate, eroding trust in high-stakes decisions [19]. Moreover, age-conditioned AHI thresholds, growth-related anatomical changes, and developmental stages are rarely embedded formally, which impairs model calibration across a child's trajectory [20]. These gaps point to the need for an integrative, semantically grounded approach that unifies evidence, preserves provenance, promotes governance, and enables transparent reasoning. Use of LLMs alone exposes the system to unnecessary risks of hallucinations. LLMs, however, provide the structure for an excellent user interface.

A knowledge graph (KG) framework offers a coherent solution by representing pediatric OSA as a network of entities and relations grounded in clinical standards and enriched with temporal structure [21]. In such a KG, concepts such as apnea subtypes, age-specific severity thresholds, symptom clusters, comorbidities, interventions, and outcomes are encoded as typed nodes, linked by clinically meaningful relations (e.g., "has-risk-factor," "treated-by," "observed-in," "responds-to"). Data from PSG, EHR, clinical notes, questionnaires, and wearables are mapped into this semantic scaffold with explicit timestamps, confidence scores, and provenance [22]. Because the KG aligns to ontologies such as SNOMED CT, LOINC, RxNorm, and HL7 FHIR resource models, it can interoperate with clinical systems, support federated analytics across institutions, and allow fine-grained validation via shape constraints [23]. Crucially, the graph becomes not just a repository but a reasoning substrate: pathway queries can retrieve evidence chains from symptom to intervention to outcome, graph algorithms can identify similar patient trajectories, and temporal rules can formalize definitions like "residual OSA six months after adenotonsillectomy" [24,25].

The conceptual architecture we advance has four pillars. First, a multi-granular semantic layer captures anatomical, physiological, behavioral, and demographic features relevant to pediatric OSA, enabling phenotype definitions that respect age and developmental stage. Second, a multimodal data integration pipeline ingests structured PSG features (AHI components, arousals, oxygen nadirs, REM-specific indices), EHR elements (ICD/SNOMED problems, medications with RxNorm codes, vitals and anthropometrics with LOINC/UCUM units), unstructured narratives from sleep reports and clinic notes, standardized questionnaire responses, and home-context wearable signals [2]. Third, a temporal model binds discrete encounters into episodes: screening, diagnostic testing, intervention, and follow-up, with start/stop times and outcome windows supporting longitudinal analyses of recurrence, growth-related changes, and treatment durability. Fourth, an interoperability layer aligns graph entities and value sets with SNOMED and FHIR, enforces data-quality constraints, and exposes standards-based APIs for clinical decision support and research [26].

Building on this architecture, we critically review the current landscape and its gaps. PSG algorithms excel at event-level accuracy but rarely encode why patterns matter clinically, for example, by connecting REM-predominant obstruction to neurocognitive outcomes or surgical decision-making [27]. EHR risk models are often trained on limited, single-center cohorts and lack explicit semantics for pediatric nuances, leading to drift when deployed elsewhere [28]. Questionnaire-based tools can be valuable triage instruments but need integration with objective measures and demographic context to avoid misclassification [29]. LLM-based summarization and question answering show promise but require grounding to verifiable facts to prevent hallucinations. A KG-centric approach addresses each shortcoming by (a) integrating modalities under shared semantics, (b) preserving audit trails from raw data to derived assertions, (c) enabling explainable paths that clinicians can inspect, and (d) supporting pediatric-specific rules that adjust as a child matures [24].

We further propose using cases and a roadmap for future research that operationalizes the KG vision [30]. In clinical decision support, the graph can power personalized screening by combining questionnaire scores with EHR comorbidities and growth metrics, flagging high-risk children even when PSG capacity is constrained [31]. For peri-operative care, temporal graph rules can estimate the risk of residual OSA after adenotonsillectomy as a function of baseline AHI, obesity status, and craniofacial findings, guiding follow-up testing [32]. In longitudinal management, the KG can monitor trajectories, linking medication adherence, allergy control, and weight change to symptom resolution or recurrence, and prompt timely interventions [33]. For research, the KG enables reproducible cohort identification (e.g., "children under 10 with REM-predominant severe OSA and asthma who received CPAP"), supports hypothesis generation around comorbidity clusters and environmental exposures, and facilitates multi-site studies through ontology-aligned data sharing. Finally, when coupled with LLMs through retrieval-augmented generation constrained by graph facts and shape validations, the system can produce trustworthy patient summaries and evidence-linked answers to clinician questions. At the same time, feedback loops convert clinician edits into ontology refinements and improvements to extraction.

In summary, pediatric OSA is prevalent, impactful, and systematically under-detected due to heterogeneous presentations and fragmented data. AI has advanced individual components such as signal processing, risk modeling, and narrative understanding, but lacks the semantic coherence, generalizability, and explainability required for dependable clinical adoption. A knowledge graph framework unifies modalities within a pediatric-aware semantic layer, encodes temporal disease courses, and anchors LLM-enabled services to verifiable facts. The contributions of this paper are threefold: it articulates a conceptual architecture for a pediatric OSA KG; it offers a critical appraisal of current AI components and

the gaps that hinder their clinical utility; and it outlines concrete use cases and a research roadmap to translate this framework into real-world decision support and multicenter research infrastructure. Through collaborative development with clinicians, informaticians, and standards bodies, this approach can help deliver earlier diagnosis, more precise treatment selection, and equitable access to high-quality pediatric sleep care.

With iterative refinement utilizing large datasets, the KG approach is expected to outperform current state-of-the-art multimodal AI models in clinical utility, explainability, and reliability because it is based on a structured, symbolic, and evidence-based system. Current AI models excel at detecting patterns across both text and images, but they frequently struggle with medical reasoning, hallucinations, explainability, and transparency. KGs address these problems by anchoring the platform in verified medical knowledge.

2. Background and Related Work

2.1. Pediatric Obstructive Sleep Apnea

Pediatric obstructive sleep apnea (OSA) arises from a confluence of anatomic, physiologic, developmental, and environmental factors that narrow or destabilize the upper airway during sleep. Etiologically, adenotonsillar hypertrophy is the dominant driver in early childhood, often interacting with craniofacial phenotypes such as midfacial hypoplasia, retrognathia, high-arched palate, and dental crowding, which reduce retropalatal or retroglossal airway caliber [34,35]. In later childhood and adolescence, obesity contributes via fat deposition around the pharyngeal airway and altered ventilatory control; neuromuscular conditions that diminish pharyngeal dilator tone further increase collapsibility [36–38]. Allergic rhinitis and chronic nasal obstruction increase inspiratory resistance and promote mouth breathing, which can worsen airway patency and disrupt sleep. These mechanisms are modulated by maturation of sleep architecture, circadian influences on ventilatory control, and growth-related changes in craniofacial structures, yielding a heterogeneous set of phenotypes that evolve across developmental stages [39,40].

Symptoms reflect this heterogeneity and often diverge from those of adults. At night, caregivers may report loud snoring, witnessed apneas or gasps, labored or paradoxical breathing, restless sleep with frequent position shifts, diaphoresis, and enuresis [41]. The presentation is highly variable, with all, a few, or even none of these symptoms. Children sometimes adopt compensatory postures (e.g., neck hyperextension) to keep the airway open [42]. Morning complaints include a sore throat and a headache. Daytime manifestations in younger children tend to involve behavioral and cognitive sequelae, such as hyperactivity, irritability, inattention, executive dysfunction, and poor school performance, rather than frank sleepiness [43,44]. By contrast, adolescents (particularly with obesity) more often exhibit classic excessive daytime sleepiness. Because these manifestations are nonspecific and vary with age, pediatric OSA is frequently underrecognized without systematic screening [2].

Comorbid conditions are common and clinically important. Asthma and allergic rhinitis can exacerbate upper-airway inflammation and resistance, creating bidirectional symptom amplification [45,46]. Hypertension and early vascular dysfunction have been observed in children with moderate-to-severe OSA, likely mediated by intermittent hypoxia, sympathetic activation, and endothelial injury [47]. Obesity not only elevates OSA risk but also attenuates responses to first-line therapies and increases the likelihood of residual disease after surgery [48]. Neurobehavioral conditions, including attention-deficit/hyperactivity disorder (ADHD), may be mimicked or worsened by sleep-disordered breathing; conversely, treating OSA can improve attention and behavior in some children [49]. Recurrent tonsillitis often coexists with adenotonsillar hypertrophy and may influence the timing

and rationale for adenotonsillectomy [50,51]. Recognizing and actively managing these comorbidities is integral to effective OSA care.

Diagnostic pathways center on objective sleep testing supported by clinical examination and standardized questionnaires. In-laboratory polysomnography (PSG) remains the reference standard in pediatrics, providing event-level annotation (apneas, hypopneas, arousals), oxygen saturation dynamics, carbon dioxide measures, sleep staging, body position, and limb movements [52]. Pediatric scoring rules differ from adult criteria and must be applied with attention to age-specific definitions and AHI thresholds. By contrast, home sleep apnea tests are generally not indicated for children because of reduced sensitivity for hypoventilation, challenges with sensor adherence and artifact in smaller patients, and the need for comprehensive physiologic monitoring to guide therapy [53]. The physical examination often reveals signs of chronic upper-airway obstruction and, in some cases, neuromuscular contributors to airway weakness. Nasal evaluations may show turbinate hypertrophy, septal deviation, or allergic stigmata; the oropharyngeal exam may reveal enlarged tonsils, elongated soft palate, or a high-arched palate; neck assessment may identify truncal/neck adiposity or masses; craniofacial assessment considers mandibular/maxillary relationships, midface hypoplasia, dental crowding, and bite class. Mouth breathing, adenoidal facies, and growth parameters provide additional context. Standardized questionnaires complement objective testing; the Epworth Sleepiness Scale for Children and Adolescents (ESS-CHAD) [54] and instruments such as the Pediatric Sleep Questionnaire (PSQ) [13] quantify symptom burden and can prioritize referrals, though they do not replace PSG.

Treatment is individualized and longitudinal. Adenotonsillectomy is first-line therapy for most children with adenotonsillar hypertrophy and clinically significant OSA, but residual disease is common in those with obesity, craniofacial anomalies, or severe baseline OSA. Continuous positive airway pressure (CPAP) is indicated when surgery is contraindicated, declined, or ineffective; its success hinges on careful mask fitting, desensitization, behavioral support, and ongoing adherence monitoring [55]. Oral appliances are not typically used in pediatric OSA because of ongoing craniofacial growth, though select orthodontic interventions (e.g., rapid maxillary expansion in appropriately selected phenotypes) may have a role within multidisciplinary care [56]. Weight management, allergy control with intranasal steroids or leukotriene modifiers, and positional therapy can be useful in specific contexts [57]. Regardless of initial therapy, structured follow-up is essential to reassess symptoms, monitor growth and weight trajectories, and repeat PSG when residual disease is suspected or clinical status changes [58].

2.2. Existing AI Approaches

Artificial intelligence has been deployed across the pediatric OSA pathway to improve detection, characterization, and care coordination, with the densest activity in PSG signal processing. Convolutional neural networks (CNNs) and temporal models (e.g., bidirectional LSTMs, temporal convolutional networks, transformers) have been trained on multi-channel PSG to automate sleep staging and event detection, reduce inter-scorer variability, and surface patterns that are challenging to quantify manually (REM-predominant obstruction, positional dependence, hypoventilation signatures) [59,60]. Self-supervised and anomaly-detection approaches leverage large volumes of unlabeled signals to pretrain robust feature extractors that adapt to pediatric physiology and smaller labeled datasets [61]. Beyond signals, clinical risk models ingest EHR-derived features and questionnaire scores to prioritize referrals, predict PSG positivity, estimate residual OSA risk after adenotonsillectomy, or recommend perioperative monitoring [20,62,63]. Feature selection techniques (e.g., LASSO, mutual information, Boruta) and ensemble learners (gradient

boosting, random forests, stacking) are commonly used, with calibration methods to improve clinical interpretability [64]. More recently, large language models (LLMs) have been piloted to summarize narrative sleep reports and clinic notes, harmonize impressions across providers, extract structured entities and relations from unstructured text, and support question answering for clinicians and families.

Despite promising results, key limitations persist. Most systems are confined to a single modality (signals, tabular EHR, or text) and a narrow task, which restricts generalizability across institutions and patient subgroups [65]. Pediatric-specific semantics, such as age-conditioned AHI thresholds, growth-related anatomy, and developmentally appropriate symptom constructs, are rarely encoded explicitly, thereby impairing calibration across developmental stages. Label scarcity and domain shift (different sensors, scoring rules, and documentation practices) challenge model robustness; class imbalance (rarer severe phenotypes) and covariate shift (obesity prevalence, referral patterns) further degrade transportability [66]. Explainability is often limited to post hoc feature attributions that do not articulate causal or mechanistic pathways clinicians can validate [19]. Finally, provenance tracking is inconsistent, complicating audit and regulatory evaluation [67]. These gaps motivate integrative approaches that unify modalities, embed pediatric semantics, and provide transparent, provenance-aware reasoning.

2.3. Knowledge Graphs in Clinical AI

Knowledge graphs (KGs) offer a principled foundation for such integration by representing clinical entities and their relationships in a computable, standards-aligned form. Examples of widely used clinical resources include SNOMED CT for problems and findings, the Unified Medical Language System (UMLS) for cross-terminology mappings, and the Observational Health Data Sciences and Informatics (OHDSI) common data model, which demonstrate how semantic harmonization enables multi-institutional analyses and reproducible phenotyping [68]. Research-oriented graphs such as MIMIC-KG illustrate how time-stamped observations, interventions, and outcomes can be linked to support cohort discovery, outcome modeling, and hypothesis generation [69]. The strengths of KG-based approaches include semantic integration across heterogeneous sources; explicit provenance and versioning that trace each assertion to its data source, timestamp, and transformation; and powerful query ability via graph languages (SPARQL/Cypher) and graph algorithms (similarity search, community detection, path reasoning).

When coupled to machine learning, graphs can inform feature construction (graph embeddings), support label-efficient learning (distant supervision), and constrain generation in LLM-based systems through retrieval-augmented pipelines that ground outputs in verifiable facts.

However, important gaps remain for pediatric OSA. Pediatric-specific ontologies are sparse, and general terminologies often lack the age-conditioned definitions and developmental nuances central to pediatric sleep medicine [70]. Multimodal integration is limited: few graphs natively integrate high-frequency physiologic signals (PSG), tabular EHR data, questionnaires, and narrative text with consistent temporal alignment and high-quality metadata [71]. Temporal reasoning itself is underdeveloped. Most graphs represent time as static attributes rather than as first-class structures that capture episodes, intervals, and event orderings necessary for trajectory-aware inference (e.g., defining residual OSA within a specific post-operative window) [72]. Addressing these gaps requires a pediatric-first KG that encodes age-specific severity thresholds and symptom constructs; ingests signals, tabular data, and narratives with synchronized timelines and uncertainty measures; and enforces provenance and validation constraints. Such a resource can serve as the seman-

tic backbone for explainable, portable, and clinically trustworthy AI systems across the pediatric OSA continuum [73].

3. Materials and Methods

3.1. A Knowledge Graph Framework for Pediatric OSA

This section lays out a practical, implementation-ready framework for a pediatric OSA knowledge graph (KG). The objective is to fuse heterogeneous clinical and contextual data into a semantically coherent substrate that supports transparent decision support, reproducible research, and longitudinal quality improvement [21]. The design emphasizes four pillars: multi-granularity, multimodality, temporality, and interoperability, which are implemented through a layered architecture comprising an ontological layer, a data layer, and LLM/NLP services bound to the KG by semantic constraints [74]. Figure 1 provides an overview of this architecture, illustrating how multimodal pediatric OSA data are harmonized into a standards-aligned knowledge graph and linked to constrained LLM/NLP services for explainable downstream use.

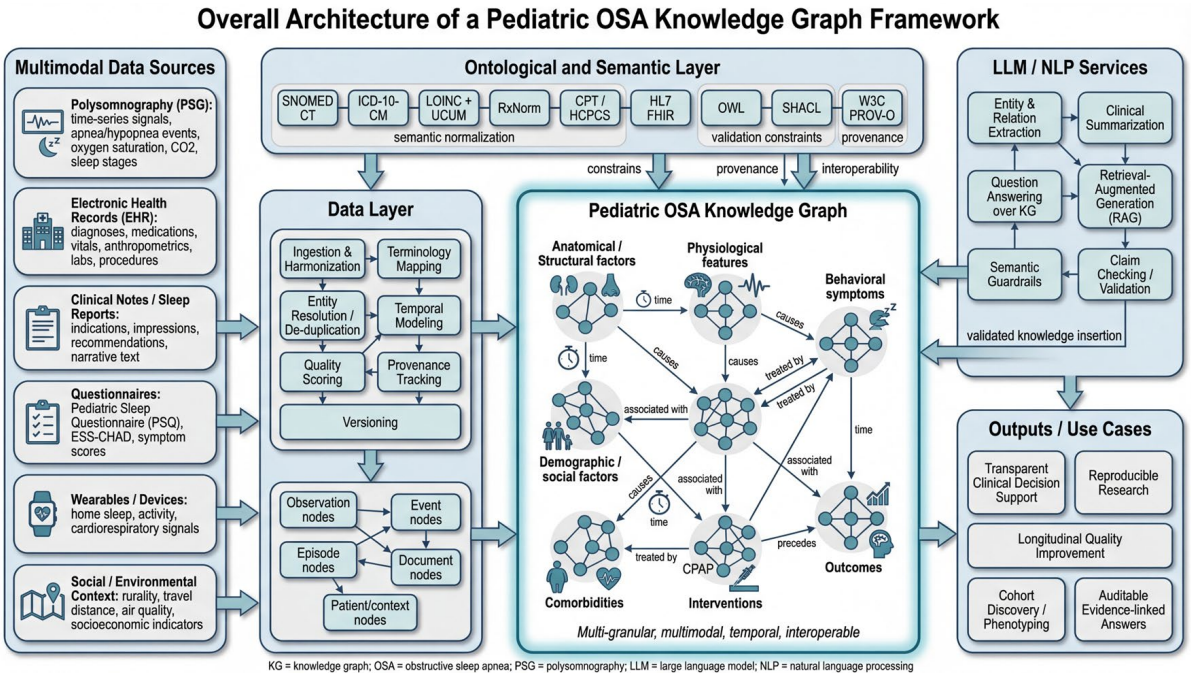


Figure 1. Layered architecture of the pediatric OSA knowledge graph framework, showing how multimodal data are harmonized through ontology-aligned pipelines and connected to LLM/NLP services under semantic constraints.

3.2. Core Design Principles

Multi-granular. Pediatric OSA emerges from interactions across anatomical structure, physiology, behavior, and demographics [75]. The KG therefore represents entities and attributes at multiple levels of abstraction [75]. At the anatomical/structural tier, classes capture craniofacial morphology (e.g., mandibular retrognathia, midface hypoplasia), adenotonsillar size, nasal patency, and airway caliber [76]. This anatomic layer can be supported by vision AI with automated analysis of the airway and facial structures. The physiological tier stores PSG-derived measures such as AHI and its components (apnea index, hypopnea index), arousal index, oxygen nadir, end-tidal/transcutaneous CO_2, sleep stage architecture, respiratory effort surrogates, and periodicity of events. The behavioral tier includes snoring frequency, witnessed apneas, parasomnias, bedwetting, and teacher/parent-reported attention or hyperactivity concerns [77]. The demographic

tier encodes age (gestational and chronological), sex, growth parameters, race/ethnicity (when appropriate and consented), and social determinants (insurance type, language, distance to specialty care, and other barriers to care). Multi-granularity allows the graph to support both fine-grained mechanistic questions (e.g., how end-tidal CO_2 trajectories relate to arousals) and higher-level clinical reasoning (e.g., whether adenotonsillectomy resolves behavioral symptoms in specific phenotypes) [38,78].

Multimodal. The framework integrates polysomnography (PSG), electronic health records (EHR), clinical notes, validated questionnaires, and wearable/device data. PSG contributes time-series events and aggregated features; EHR contributes diagnoses, medications, vitals, anthropometrics, laboratory results, and procedures; narrative notes add nuanced clinical impressions and rationales; questionnaires such as the Pediatric Sleep Questionnaire (PSQ) and ESS-CHAD encode standardized symptom burdens; and wearables provide home-context sleep/wake and cardiorespiratory signals [52,53]. Modality-specific uncertainty is preserved as first-class metadata to avoid over-confidence when modalities disagree [79].

Temporal. Pediatric OSA trajectories evolve with growth, weight change, surgery, orthodontic development, and comorbidity management [79]. The KG models time explicitly: clinical "episodes" (screening → diagnostic PSG → intervention → follow-up → recurrence) are linked by temporal relations (precedes/follows, overlaps, during). Observations carry effective time ranges; interventions carry start/stop times and dose/exposure details; outcomes are measured at defined windows (e.g., 6-week, 6-month, 12-month follow-ups) [52]. Figure 2 illustrates this temporal patient journey and shows how multimodal observations, interventions, and outcomes are anchored to longitudinal episodes rather than treated as isolated clinical snapshots. This enables longitudinal queries (e.g., residual OSA six months after adenotonsillectomy in children with obesity) and supports causal modeling and survival analysis layered atop the graph.

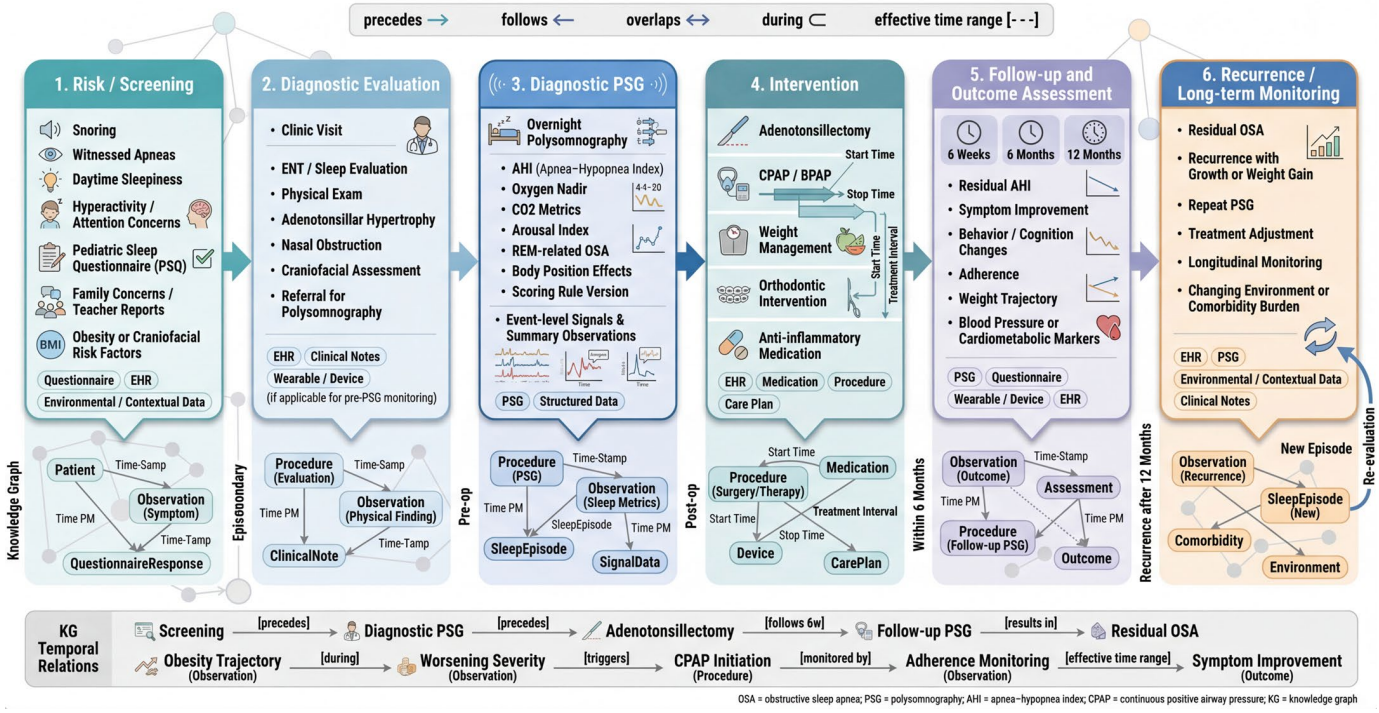


Figure 2. Temporal modeling of pediatric OSA episodes across screening, diagnosis, treatment, follow-up, and recurrence, enabling longitudinal reasoning over interventions and outcomes.

Interoperable. To ensure portability and safety, entities are mapped to clinical standards. Conditions, findings, and procedures align to SNOMED CT; diagnoses to ICD-10-CM (with pediatric-appropriate use); observations to LOINC with UCUM units; medications to RxNorm; and care processes to CPT/HCPCS where relevant [80]. Data exchange is mediated via HL7 FHIR resources (e.g., Patient, Encounter, Observation, Condition, Procedure, Medication Statement, Device, QuestionnaireResponse). Interoperability allows federated learning across institutions and simplifies regulatory review, audit, and integration with clinical systems [81].

3.3. Ontological Layer

The ontology provides the semantic backbone. **Concepts** include apnea types (obstructive, central, mixed, hypopnea per scoring rule), severity categories (age-appropriate AHI thresholds), symptom clusters (nocturnal vs. daytime behavioral/cognitive), comorbidities (obesity, asthma, allergic rhinitis, craniofacial anomaly, neuromuscular disease, hypertension, ADHD), and interventions (adenotonsillectomy, CPAP/BPAP, weight management, orthodontic therapies, anti-inflammatory regimens) [14]. Severity classes are defined with pediatric-specific cut points and may differ for adolescents approaching adult criteria; these distinctions are encoded as age-conditioned class axioms rather than static labels [82].

Relationships capture clinically meaningful semantics, including *has-risk-factor* (e.g., obesity → OSA) [52], *treated-by* (severe OSA → CPAP) [83], *observed-in* (hypoventilation → REM sleep), *responds-to* (behavioral symptoms → adenotonsillectomy), *contraindicated-with* (specific craniofacial anomalies → oral appliances), *co-occurs-with* (OSA ↔ asthma), *worsened-by* (upper respiratory infection → OSA severity) [84], and *precedes/follows* (pre-op PSG → surgery → post-op PSG). Temporalized relations allow reasoning like "post-operative AHI < 2 within 6 months [85]".

The Shapes Constraint Language (SHACL) is a language for validating Resource Description Framework (RDF) graphs against a set of conditions. SHACL has a core specification, extensions for related query languages, User Interface generation, rule-based reasoning and a data profiling mechanism. SHACL models are defined in terms of constraints on the content, structure, and meaning of a graph. Wherever possible, the ontology **leverages existing pediatric vocabularies and meta-thesauri** (e.g., SNOMED CT hierarchies for pediatric findings, UMLS mappings, FHIR ValueSets for pediatric vitals) and extends them minimally to cover OSA-specific constructs not well represented (e.g., child-specific scoring variants, questionnaire score semantics) [52,86]. Constraints are formalized using OWL restrictions for taxonomic reasoning and SHACL shapes for data validation (e.g., a *PSGObservation* must reference a LOINC code and UCUM unit, and an *AHI* value must be associated with a scoring rule version) [79]. Provenance is modeled with W3C PROV-O so that every edge and attribute can be traced to its data source, timestamp, and transformation step, enabling auditability and trust [13].

The ontology layer manages the transition from infant to pediatric to adult sleep physiology and associated diagnostic assessments via a structured approach that maps dynamic developmental changes such as the definition of AHI which varies by age into a structured framework. To ensure clinical utility and explainability, this involves moving from static, stage-specific ontologies to a "developmental continuum" model that uses shared semantic anchors (e.g., SNOMED-CT, LOINC) to bridge pediatric and adult data systems.

3.4. Data Layer

The data layer operationalizes ontology-aligned ingestion, harmonization, and storage.

Polysomnography (structured features). The pipeline extracts per-study aggregates (AHI and sub-indices, arousal index, oxygen nadir, REM-specific indices, CO_2 metrics),

per-event sequences (start/stop, type, stage, body position), and channel-level quality indicators. Signal-processing metadata (scoring rule version, scorer identity, automated vs. manual adjudication) is preserved. Features are stored both as time-stamped events linked to *SleepEpisode* and as summary *Observation* nodes with LOINC/UCUM annotations, enabling analyses from both event-level and study-level perspectives [78].

EHR (ICD, medication, vitals, demographics). Diagnoses (e.g., OSA, obesity, asthma), procedures (adenotonsillectomy, orthodontic interventions), medications (intranasal steroids, leukotriene antagonists), anthropometrics (height, weight, BMI z-scores), and vitals (blood pressure percentiles) are normalized to SNOMED/ICD/RxNorm/LOINC and linked to encounters and episodes. Growth trajectories are modeled as sequences to support phenotype definitions that depend on temporal change (e.g., a rapid increase in BMI percentile preceding OSA onset).

Sleep study reports (unstructured). Narrative sections ("Indications," "Impression," "Recommendations") often contain nuanced rationales and clinical caveats. Reports are stored as documents with stable identifiers; LLM/NLP services extract entities (e.g., "REM-predominant OSA," "positional dependence") and relations (e.g., "surgery recommended because. . .") into the KG with confidence scores and provenance backpointers to text spans, enabling human verification [52,53].

Survey/questionnaire data (e.g., PSQ). Raw item responses and total/subscale scores are modeled explicitly, with score interpretations (e.g., cutoffs indicating high risk) encoded as computable rules. This allows queries such as "children with high PSQ score but normal in-lab AHI" to investigate discordant cases [79].

Social/environmental data (e.g., rurality, pollution). Geocoded features (appropriately de-identified and privacy-protected) capture rural–urban classification, travel distance/time to sleep centers, local air quality indices, secondhand smoke exposure proxies, and socioeconomic indicators. These are linked to patients via time-varying residence nodes so environmental changes (e.g., seasonality in pollution) can be studied in relation to symptoms and PSG metrics [87].

Architecture and storage. A hybrid approach supports both semantic reasoning and high-performance analytics: an RDF/OWL triple store with SPARQL endpoints provides standards-based semantics and reasoning, while a property-graph mirror (e.g., Neo4j) accelerates path queries and graph algorithms (similar patient retrieval, community detection). ETL pipelines enforce SHACL constraints, perform de-duplication and entity resolution, and attach data quality scores (completeness, plausibility, consistency). Versioning enables "time-travel" analyses and reproducibility (e.g., rerunning a cohort definition against the KG as of a specific date) [52,79].

3.5. LLM and NLP Integration

LLMs and classical NLP are coupled with the KG to extract knowledge, summarize context, and support question answering, always under **graph-aware constraints** to preserve clinical fidelity.

Entity and relation extraction. A hybrid pipeline combines dictionary/ontology lookups, weakly supervised sequence labeling, and LLM few-shot prompting to extract entities (diagnoses, symptoms, measurements, recommendations) and relations ("treated-by," "observed-in," "worsened-by") from sleep reports, clinic notes, and literature. Extracted candidates are normalized to controlled vocabularies, disambiguated using KG context (e.g., distinguishing "central apnea" as an event vs. a diagnosis), and inserted only when SHACL constraints pass. Each assertion carries confidence, provenance (document, section, sentence offsets), and model version, enabling selective human review for low-confidence triples [88].

Clinical summarization. For point-of-care use, LLMs generate concise, structured summaries of a patient's OSA trajectory (indications, PSG findings, comorbidities, prior interventions, outcomes). Summaries are composed by retrieving relevant subgraphs (patient $\rightarrow$ episodes $\rightarrow$ observations/interventions) and rendering them via controlled templates augmented by LLM natural-language fluency [83]. Consistent with the temporal patient journey shown in Figure 2, these summaries are organized around linked longitudinal episodes rather than isolated encounters, allowing clinicians to interpret findings in the context of disease progression, treatment response, and recurrence. KG facts are canonical and cannot be altered by the generator; where uncertainty exists (e.g., conflicting scores), the summary explicitly flags it and links to source nodes for drill-down inspection by the clinician [53].

Question answering over the KG. User queries are translated into graph operations through a retrieval-augmented generation (RAG) layer. First, the system retrieves a minimal subgraph satisfying the information need (e.g., "children < 10 years with severe OSA who had adenotonsillectomy and residual AHI $\geq$ 5 within 12 months"); second, SPARQL/Cypher executes the query; third, the LLM verbalizes results with explanation and citations to graph nodes/edges. "Semantic guardrails" constrain the LLM only to reference retrieved graph facts; responses include evidence links and uncertainty qualifiers [79]. For multi-site deployments, the RAG layer supports privacy-preserving federation by shipping queries to local KGs and aggregating de-identified results.

KG-constrained accuracy. To minimize hallucinations and support regulatory-grade reliability, LLM/NLP components are tightly governed by the pediatric OSA knowledge graph and its semantic control layer throughout extraction, summarization, and question answering. As illustrated in Figure 3, the workflow begins with graph-aware retrieval and knowledge extraction, in which candidate entities and relations from notes, reports, questionnaires, and other inputs are normalized to controlled vocabularies, disambiguated using KG context, and treated as candidate assertions rather than trusted facts. Generation is bound to KG constraints at three checkpoints: (1) pre-generation retrieval, which limits the model's context to verified graph facts and a minimal evidence-bearing subgraph [89]; (2) **during-generation validation**, which blocks unsupported claims and prevents outputs that would contradict graph state, semantic mappings, or SHACL constraints; and (3) **post-generation claim checking**, which re-parses the generated text, reconciles each statement against the KG, and emits discrepancy reports or uncertainty flags when inconsistencies are detected [90]. Only validated assertions are inserted back into the KG with provenance, confidence scores, and model-version metadata, ensuring that the graph remains the canonical source of truth. Continuous learning loops capture clinician feedback on extracted triples, summaries, and answers, converting corrections into training data and ontology updates.

Equity Auditing and Bias Mitigation. To ensure that the proposed KG framework supports equitable clinical decision-making, we incorporated a dedicated equity auditing and bias mitigation layer into the platform. Demographic and social determinant variables (e.g., race/ethnicity, insurance status, language, and geographic access to care) are explicitly represented as nodes and attributes within the KG, enabling subgroup-level analyses. Model outputs derived from the KG are evaluated using fairness-aware metrics, including subgroup-specific sensitivity, specificity, calibration, and disparity measures. To enforce structural and semantic integrity, we apply SHACL validation rules that detect imbalances in data representation, flag disproportionate associations between clinical recommendations and demographic attributes, and ensure that inferred relationships meet predefined equity constraints. In addition, the KG's built-in provenance model enables traceability of each inference to its underlying data sources, supporting transparent bias auditing. When disparities are identified, mitigation strategies, including reweighting, stratified resam-

pling, and ontology refinement, are applied iteratively, with clinician (human-in-the-loop) oversight, to reduce bias and improve model fairness. This approach ensures that the KG framework not only integrates multimodal data but also actively monitors and allows for mitigation of systemic inequities present in historical EHR data and other databases.

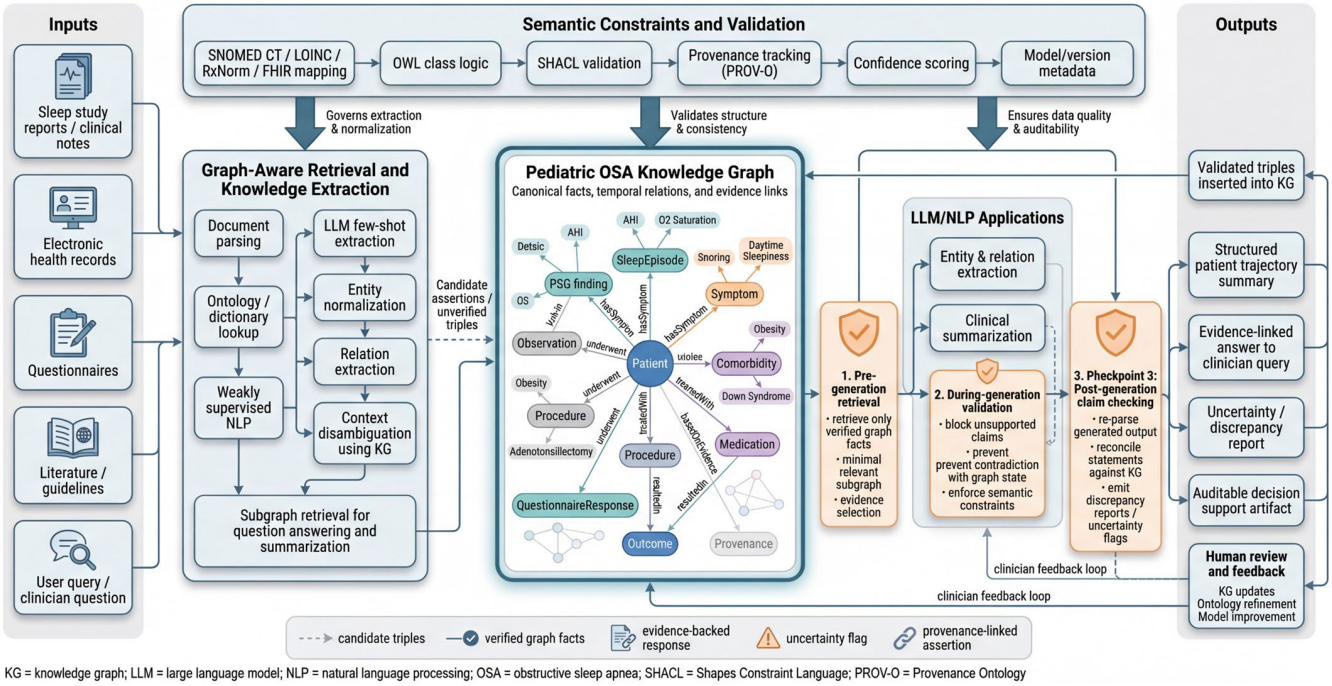


Figure 3. Knowledge-graph-constrained LLM/NLP workflow for extraction, summarization, and question answering with semantic guardrails and auditability.

Validation Roadmap. The validation strategy for AI analysis of obstructive sleep apnea (OSA) involves several key steps to ensure the accuracy and reliability of the AI models. Accurate and reliable Training Datasets are of paramount importance. The large internal data repositories available to this group will be leveraged as well as open-access datasets such as the ISRUC-Sleep from the Sleep Medicine Centre of the Hospital of Coimbra University. Clinical validation, conducted in the clinical setting, is planned for future trials in the University Medical Center Sleep Medicine Clinics. The quality and completeness of the training and validation datasets available from the internal datasets has already been assessed.

Model performance. Transparency, explainability, provenance, and governance metrics for the datasets are available. These datasets are fully anonymized to preserve patient privacy with supervision by the Institutional Review Board.

Summary. The proposed KG framework transforms pediatric OSA data into a computable, explainable, and standards-aligned representation that reflects the condition's developmental and multimodal complexity [21]. The patient journey is summarized in Figure 4. Multi-granularity supports rich phenotype definitions; multimodality preserves complementary evidence; temporal modeling enables trajectory-aware reasoning; and interoperability ensures integration with clinical systems and research networks [72,81,91]. An ontology grounded in pediatric practice and reinforced by OWL/SHACL provides semantic rigor, while the data layer delivers robust pipelines for heterogeneous sources, quality control, and analytics [92]. Finally, LLM/NLP components tethered to the KG through RAG and semantic guardrails unlock scalable extraction, trustworthy summarization, and auditable question answering, positioning the framework for real-world clinical decision support and multicenter research [90,93].

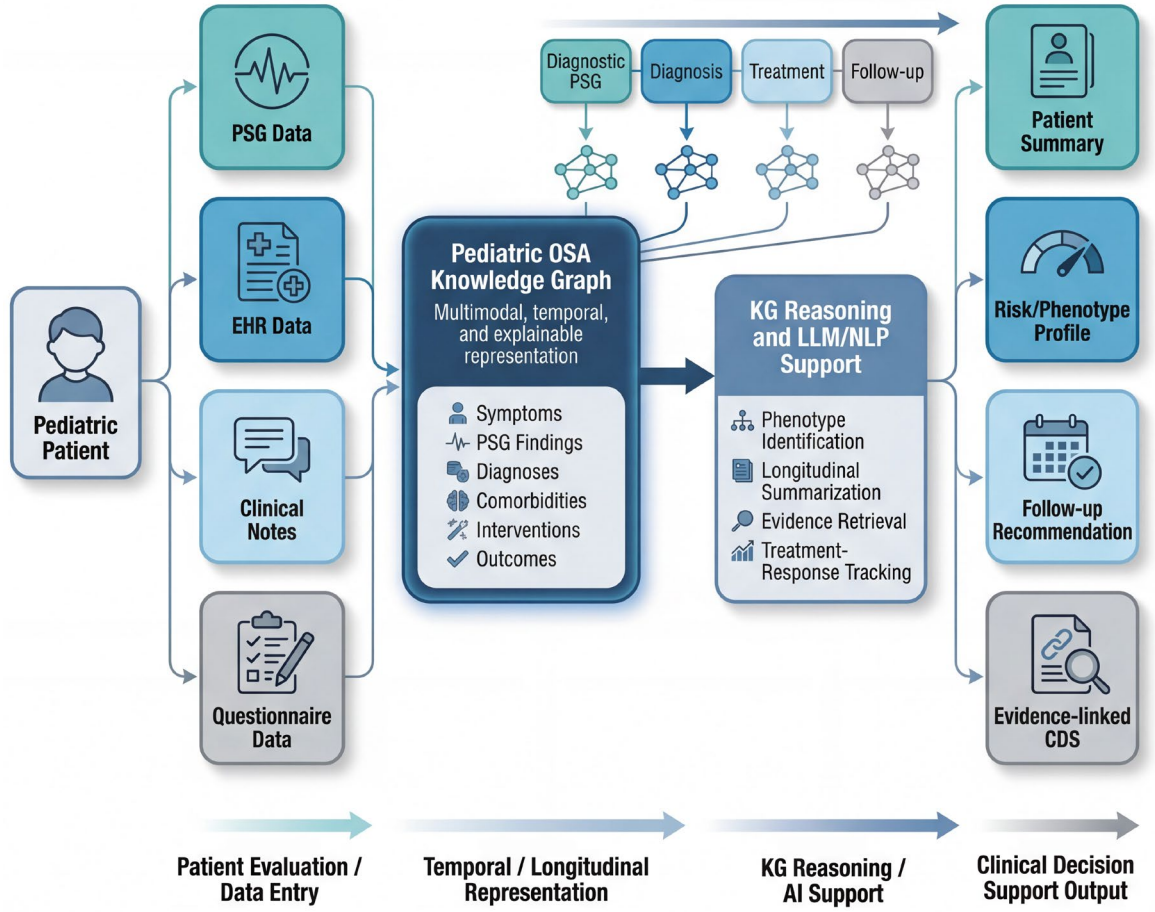


Figure 4. High-level flow of a single pediatric OSA patient journey from multimodal data capture to knowledge graph-enabled clinical decision support.

4. Results

4.1. Use Cases and Applications

The pediatric OSA knowledge graph (KG) enables a family of applications that translate heterogeneous data into timely, explainable decisions and reproducible science [94]. By encoding symptoms, measurements, diagnoses, interventions, and outcomes as typed nodes with clinically meaningful relations and timestamps, the KG becomes both a unifying data substrate and a reasoning layer. As this is a proposed conceptual framework without empirical implementation to date, the following use cases are illustrative of the framework's intended capabilities and are grounded in analogous work in the broader clinical informatics literature. Below, we detail how this foundation powers clinical decision support, accelerates research and discovery, and enables interpretable AI models designed for pediatric care [20,66].

4.2. Clinical Decision Support

In personalized screening, the KG fuses questionnaire responses with EHR context to estimate a child's likelihood of clinically significant OSA before polysomnography. Questionnaire items (for example, PSQ subscales or ESS-CHAD scores) are represented as first-class observations linked to encounters, while the EHR contributes demographics, anthropometrics (BMI-for-age z-scores), comorbidities (asthma, allergic rhinitis, ADHD), medications (intranasal steroids, leukotriene modifiers), prior procedures, and social determinants (distance to a sleep center, insurance type) [15,57,95]. Graph queries assemble a patient-specific subgraph; a screening model then operates on graph features rather

than flat tables, preserving age-conditioned semantics (e.g., pediatric AHI thresholds) and temporality (e.g., recent weight gain). The system returns a calibrated risk estimate, the top evidence paths supporting that estimate (such as "high PSQ snoring score + obesity + allergic rhinitis") [15], and next-step options tailored to local capacity—watchful waiting with allergy control, expedited PSG referral, or pre-test counseling for families [96]. Decision curve analysis and impact metrics (positive predictive value at fixed referral slots, missed-case rate) guide threshold selection, enabling clinics to match screening to resource constraints while maintaining equity [97,98].

For postoperative risk assessment, the KG tracks the full adenotonsillectomy episode—preoperative PSG features (baseline AHI, oxygen nadir, REM-predominant or positional dependence), perioperative factors, and longitudinal follow-up findings [99]. Temporal relations bind these to outcome windows (e.g., residual OSA within 6–12 months) [32]. A predictive model on the graph estimates residual OSA risk conditional on phenotype (obesity, craniofacial anomaly, asthma control) and recommends surveillance strategies, such as timing for repeat PSG or prioritizing CPAP titration [33,100]. The output is accompanied by explanation subgraphs that show how features combine "severe baseline OSA + obesity + persistent allergic rhinitis $\rightarrow$ high residual risk" and by actionable levers (optimize allergy therapy; intensify weight management; plan earlier post-op testing). Because provenance is explicit, clinicians can audit which measurements, rules, and sources contributed to each recommendation.

A sleep-specialist assistant layers case-based reasoning atop the KG. When presented with a new patient, the assistant retrieves clinically similar trajectories using graph similarity (embedding-based nearest neighbors constrained by age and key comorbidities) and surfaces outcomes of comparable cases, linked guidelines, and evidence [101]. For example, for an 8-year-old with REM-predominant OSA, obesity, and asthma, the assistant may display cohorts where adenotonsillectomy improved AHI but residual disease persisted in obese children [78,99], along with comparative outcomes for early CPAP initiation versus staged management [100]. The assistant's narrative is generated through retrieval-augmented generation constrained by graph facts, ensuring that all statements are grounded in verifiable nodes and edges and are accompanied by citations and confidence levels [102].

4.3. Research and Discovery

The KG streamlines hypothesis generation by enabling population-level patterns to be explored without bespoke data wrangling. Investigators can query age-stratified AHI distributions, examine REM-predominant patterns by sex or obesity status [10], or test associations between comorbidity clusters (asthma + allergic rhinitis) and oxygen desaturation burden [103,104]. Because temporal intervals are explicit, analyses can differentiate pre- and post-intervention trajectories and evaluate seasonality (e.g., allergy peaks) or growth-related changes [105]. Equity-focused questions such as differences in access to PSG or time-to-treatment by rurality or insurance are likewise computable, with sensitivity analyses facilitated by provenance and data-quality scores [95,98].

Cohort identification becomes reproducible and portable. Phenotypes are formalized as graph patterns—for instance, "children $<$ 10 years, baseline AHI $\geq$ 10, REM-predominant events, comorbid asthma, underwent adenotonsillectomy, residual AHI $\geq$ 5 within 12 months." These patterns can be versioned, shared, and re-executed across institutions that map their data to the same ontology, enabling multi-site studies without reinventing inclusion logic [106,107]. Graph embeddings complement rule-based phenotypes by clustering children with similar disease trajectories, helping researchers discover latent subtypes [108,109] (e.g., positional-dominant OSA with normal BMI versus obesity-driven OSA with hypoventilation features) that might merit different follow-up [110].

Augmented literature review and knowledge synthesis leverage the KG's ontology to index publications and guideline statements by the concepts they address (apnea subtype, intervention, age band, comorbidity) [111]. An LLM retrieves graph-linked evidence to answer targeted questions: "What is the expected change in AHI after adenotonsillectomy in obese children under 10?" and produces an evidence table with citations, effect sizes, and study contexts [102,112]. Contradictions across studies are flagged when graph-encoded populations or outcome definitions differ, and new papers are incrementally integrated by extracting entities and relations with provenance back to text spans [113]. The result is a living, pediatric-aware evidence graph that reduces duplication and accelerates consensus building [114].

4.4. Interpretable AI Models

Graph neural networks (GNNs) operating on the pediatric OSA KG enable outcome prediction while preserving structure and context. Heterogeneous GNNs encode node and edge types (patient, observation, questionnaire score, procedure; "has-measurement," "treated-by," "precedes"), and temporal GNN variants incorporate event timing and episode boundaries [115]. Tasks include predicting residual OSA after surgery, likelihood of CPAP adherence at three months [116], or risk of neurobehavioral impairment given baseline PSG and comorbidity profile. Training uses stratified, time-split validation to respect temporality; evaluation emphasizes AUROC/AUPRC alongside calibration (Brier score, reliability diagrams) and clinical utility (decision curves) [97,117]. Because inputs are graph-structured, models can naturally integrate multimodal evidence with pediatric-specific semantics, improving transportability relative to flat-feature baselines [115].

Interpretability is delivered through explainable pathways that connect symptoms to interventions and outcomes. Post hoc explainers tailored to graphs (e.g., GNNExplainer, PGExplainer) identify compact subgraphs including specific questionnaire items, obesity status, REM-predominant events, allergy control that were most influential for a prediction [118,119], while counterfactual analysis searches for minimal graph edits that would change the recommendation (for example, "if allergic rhinitis were controlled and BMI-z decreased by 0.5, residual-risk classification would drop below threshold"). These explanations are rendered as human-readable narratives with clickable evidence nodes, aligning with clinical reasoning and supporting shared decision-making with families [120,121]. Crucially, explanations are validated against domain rules encoded in the ontology (age-appropriate thresholds, contraindications), ensuring fidelity to pediatric practice [108]. Uncertainty estimates (e.g., conformal prediction sets) accompany outputs so clinicians can weigh predictions appropriately and prioritize confirmation testing when model confidence is low [122].

Together, these applications illustrate how a pediatric-first KG transforms fragmented data into actionable, auditable intelligence. Clinical teams gain transparent tools for screening, peri-operative planning, and longitudinal management; researchers gain reproducible cohorts, rapid hypothesis testing, and living evidence maps; and model developers gain a semantically rich substrate that supports interpretable, temporally aware learning [123]. The common thread: explicit semantics, provenance, and time ensures that advances in automation translate into trustworthy improvements in pediatric sleep health [124].

4.5. Preliminary Feasibility Demonstration

To provide initial empirical support for the proposed knowledge graph (KG) framework, we constructed a small synthetic pilot dataset (n = 24) reflecting the multimodal structure of pediatric OSA data described in this architecture, including polysomnography-derived

metrics (AHI, oxygen nadir), questionnaire-based symptom burden (PSQ), comorbidities (e.g., obesity, asthma, allergic rhinitis), and longitudinal outcomes following intervention.

To directly address clinically relevant edge cases, we explicitly incorporated discordant phenotypes, including (i) patients with high physiologic severity (AHI $\geq$ 10) but low reported symptom burden and (ii) patients with low AHI but high symptom burden. These discordant cases are commonly encountered in pediatric sleep medicine and represent a key limitation of existing single-modality models.

Within the KG, these cases were represented as structured subgraphs linking measurements, symptom scores, comorbidities, and outcomes. Figure 5 illustrates an example of a discordant phenotype in which severe AHI coexists with low symptom burden. The graph preserves this conflict while contextualizing it with additional features, including obesity, asthma, and REM-predominant events.

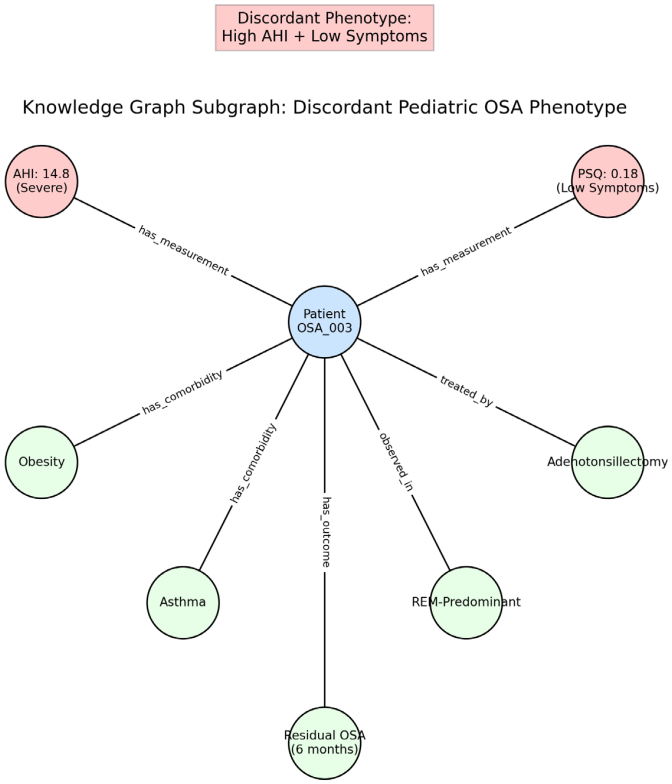


Figure 5. Knowledge graph subgraph illustrating a discordant pediatric OSA phenotype. This example demonstrates a patient with severe physiologic disease (AHI = 14.8 events/h) but low reported symptom burden (PSQ = 0.18). The knowledge graph explicitly represents conflicting clinical signals while preserving contextual relationships, including comorbidities (obesity, asthma), REM-predominant disease, intervention (adenotonsillectomy), and longitudinal outcome (residual OSA at 6 months). Nodes: SNOMED-CT entities; edges: causal relations with probabilistic weights. This structure enables downstream graph-based learning and retrieval methods to identify similar discordant cases and support explainable clinical reasoning.

To demonstrate graph-based reasoning, we applied a simple similarity-based retrieval over graph-derived feature embeddings. For a representative discordant case, the system identified clinically similar patients sharing features such as REM predominance and obesity, many of whom exhibited persistent disease following intervention. This suggests that the graph structure captures meaningful clinical relationships beyond individual variables.

In addition, we implemented a prototype retrieval-augmented generation (RAG) workflow constrained by the KG. Given a clinician-style query (e.g., "Why does this patient have severe AHI but minimal symptoms?"), the system retrieved relevant subgraph elements

and generated an explanation grounded in structured evidence (e.g., REM predominance, obesity, comorbid asthma), demonstrating the feasibility of explainable, graph-constrained reasoning. The RAG architecture requires iterative refinement, as would be expected for any advanced platform.

While this pilot is limited in scale and uses synthetic data, it provides proof-of-concept evidence that the proposed KG + GNN + RAG architecture can (i) represent multimodal pediatric OSA data, (ii) explicitly encode discordant clinical signals, and (iii) support interpretable reasoning over complex patient profiles. These results motivate future validation using real-world, multi-institutional datasets.

5. Discussion—Challenges and Future Directions

5.1. Data Availability and Quality

The most immediate barrier to a robust pediatric OSA knowledge graph (KG) and its downstream analytics is the scarcity of labeled multimodal datasets that span signals, tabular clinical variables, narratives, and longitudinal outcomes. Pediatric cohorts are small by nature, heterogeneity across developmental stages is high, and clinically adjudicated labels, specifically for nuanced constructs such as REM-predominant obstruction, hypoventilation, or behavioral outcomes, are costly to obtain and often inconsistent across centers [20]. Label noise arises from inter-scorer variability in PSG, evolving scoring manuals, incomplete documentation of comorbidities, and variation in questionnaire administration [125,126]. Domain shift further degrades model transportability: sensors and montage configurations differ across labs; EHR coding practices, care pathways, and referral biases vary across systems; and population mix (e.g., obesity prevalence, craniofacial anomalies) varies by geography and time [127]. Addressing these deficits requires deliberate data governance and engineering: standardized ingestion pipelines with explicit versioning of scoring rules; SHACL- or schema-based validation to enforce semantic and unit consistency [128]; calibrated data-quality metrics (completeness, plausibility, concordance) attached to each node and edge [129]; and targeted labeling campaigns that prioritize high-impact ambiguities identified by disagreement mining. Semi-supervised learning and weak supervision can amplify limited labeled sets, while active learning can focus expert review where model uncertainty is greatest [130]. Finally, benchmark suites (i.e., shared, de-identified task definitions withheld-out external test sets) are essential to evaluate generalization and fairness across ages, racial/ethnic groups, and comorbidity strata [98].

De-identification and data sharing present additional constraints that are particularly acute in pediatrics. Standard date shifting, pseudonymization, and removal of direct identifiers often leave residual re-identification risk when rare phenotypes, detailed timelines, or free-text notes are included [131]. Safe collaboration demands a layered strategy: rigorous de-identification with risk assessment tailored to pediatric rarity [132]; differential privacy for aggregate releases when appropriate [133]; governance frameworks with clear data-use agreements, provenance capture, and audit trails; and technical enablers such as secure data enclaves, federated learning, and federated analytics that allow models to be trained across institutional KGs without raw data leaving local control [134,135]. Synthetic data can help with software testing and education but should not be treated as statistically equivalent to source data without careful utility and privacy evaluation. For free text, structured extraction pipelines must retain bidirectional links to redacted spans so that clinicians can verify assertions without exposing protected information.

5.2. Ontology Development

Progress is gated by the scarcity of pediatric OSA-specific semantic standards. General terminology captures many building blocks (e.g., problems, findings, procedures, medica-

tions, and measurements) but rarely encodes age-conditioned definitions, developmental considerations, or pediatric scoring nuances central to sleep medicine [136]. A sustainable path forward is a layered ontology that reuses established vocabularies for core entities while adding minimal pediatric extensions expressed in OWL and validated by SHACL shapes [128]. Age-conditioned axioms should be first-class semantics rather than free-text annotations. For example, an AHI of 13 would indicate severe obstructive sleep apnea in a 3-year-old, while this same AHI would represent only mild obstructive sleep apnea in a 19-year-old. Likewise, questionnaire constructs (item-level responses, subscales, cut points) need computable definitions to support reproducible phenotyping and decision support. Each extension must be versioned, provenance-stamped, and accompanied by competency questions that demonstrate clinical utility (e.g., "identify children <10 years with REM-predominant severe OSA and asthma who have residual disease 6–12 months post-adenotonsillectomy").

Community-driven curation is essential to ensure breadth, depth, and adoption. A multi-stakeholder consortium including sleep clinicians, pediatricians, otolaryngologists, orthodontists, informaticians, and patient advocates should govern change proposals, editorial workflows, and release cadence. Transparent contribution pipelines (issue tracking, pull requests, review boards), reference implementations (public SHACL validators, example SPARQL queries), and alignment artifacts (cross-walks to SNOMED CT, LOINC, RxNorm, HL7 FHIR ValueSets) will lower integration costs [137]. Crucially, the community must commit to longitudinal stewardship: updating axioms when scoring rules evolve, adding new device metrics as technology changes, and retiring deprecated terms with clear migration paths so historical data remain queryable.

5.3. LLM Alignment and Trust

Large language models (LLMs) unlock powerful capabilities—entity/relation extraction from notes and reports, clinical summarization, and question answering—but their unconstrained use risks hallucination, subtle clinical inaccuracies, and misplaced confidence. Trustworthy deployment demands alignment strategies that go beyond prompt engineering. Instruction-tuned, pediatric-focused corpora improve baseline adherence to domain style and content, but fine-tuning alone does not guarantee factuality or provenance. A hybrid architecture should bind generation to KG facts through retrieval-augmented generation (RAG), constrained decoding, and post hoc claim checking [102]. In practice, this means pre-restricting the model's context to a verified, minimal subgraph; enforcing output schemas and ontological constraints during decoding (e.g., preventing age-inappropriate severity labels); and re-parsing generated statements to cross-verify each claim against the KG, flagging unsupported assertions for clinician review. For extraction tasks, span-level provenance and confidence scores must accompany every triple, enabling selective human validation and active-learning loops that prioritize low-confidence or high-impact edges.

Choosing among prompt engineering, fine-tuning, and KG-constrained hybrids is not binary; it is task-specific. Prompt engineering offers speed and flexibility for exploratory tasks and UI iteration but requires strong guardrails and conservative defaults. Supervised fine-tuning on curated pediatric corpora yields better consistency for repetitive tasks (e.g., report summarization) and reduces prompt brittleness. The highest assurance arises from hybrid pipelines that combine RAG with symbolic checks: the KG supplies facts and constraints; the LLM provides language fluency and abstraction; and a verifier reconciles outputs with graph state. Governance should include red-teaming with clinically plausible adversarial prompts, calibration assessment (e.g., expected calibration error, conformal prediction sets), and continuous monitoring for drift, with rollback mechanisms and model cards that disclose training data, intended use, and known failure modes.

5.4. Human-in-the-Loop

This adjunctive consultant platform is meant to allow for human-in-the-loop oversight at all stages of the evaluation and decision-making process. We strongly believe that there is no replacement for the physician's evaluation and the understanding gleaned from the patient–physician relationship and interactions. Essentially, the system would empower the physician at all stages. This information would be used to enhance the platform in ongoing iterations.

5.5. Deployment and Adoption

Even the best models fail without thoughtful integration into clinical workflows. The KG and its services should be surfaced through standards-based interfaces (e.g., SMART-on-FHIR apps, CDS Hooks for event-triggered recommendations, and FHIR Subscriptions for longitudinal monitoring) so that decision support appears in the right place, at the right time, with minimal context switching [138]. Human-factors engineering is critical: interfaces must present concise recommendations with evidence paths, uncertainty indicators, and "why this, why now" explanations, while avoiding alert fatigue via configurable thresholds and suppression rules. Implementation should follow iterative Plan-Do-Study-Act cycles with champions in sleep clinics and peri-operative teams, accompanied by training materials and feedback channels that route clinician input back to ontology and model owners. Equity must be monitored explicitly: dashboards should track performance and access across age, sex, race/ethnicity, language, rurality, and payer status, with remediation plans when disparities emerge [98].

Validation and regulatory pathways require the same rigor expected of other pediatric clinical technologies. Before-and-after or stepped-wedge designs can evaluate impact on referral appropriateness, time-to-diagnosis, residual disease detection, and patient-reported outcomes, while external validation across sites quantifies transportability. For software as a medical device, teams should adopt quality and safety frameworks (risk management, software lifecycle, usability engineering, cybersecurity) and maintain algorithm change protocols that pre-specify what can evolve without re-clearance and how performance will be re-verified. Real-world performance monitoring must be continuous: model drift detectors, audit logs for recommendations and overrides, and incident response playbooks for erroneous outputs. Documentation utilizing model cards, datasheets, validation reports, and post market surveillance summaries should be treated as living assets. Adoption hinges on delivering measurable value to clinicians and families: faster, fairer screening; clearer post-operative planning; and more reliable long-term management. By aligning technical advances with governance, safety, and usability, the pediatric OSA KG can translate from promising research into dependable, equitable care.

6. Conclusions

Pediatric obstructive sleep apnea (OSA) demands tools that can reconcile heterogeneous evidence, honor developmental nuance, and produce recommendations clinicians can trust. A knowledge graph (KG)-based framework provides that foundation by encoding pediatric OSA as a semantically rigorous, temporally aware, and interoperable representation that integrates signals (PSG), clinical data (EHR), narrative reports, questionnaires, and context such as environment and growth. Anchoring analytics to this graph enables multimodal reasoning, preserves provenance, and supports transparent explanations—transforming AI from narrow, siloed point solutions into a cohesive, explainable system of care. By coupling the KG with graph-aware learning (e.g., GNNs) and KG-constrained generation (RAG + ontological checks), the approach delivers predictions and summaries that are not only accurate, but auditable and aligned with pediatric practice.

Realizing this vision requires a coordinated, open, and sustained collaboration. We call for a multi-stakeholder consortium and standards bodies to co-develop: (i) a pediatric OSA ontology with age-conditioned axioms and SHACL validation; (ii) reference ETL pipelines and federated governance that enable privacy-preserving, multi-site data integration; (iii) benchmark tasks and external test sets spanning diverse populations; and (iv) clinical decision support pilots embedded via SMART-on-FHIR/CDS Hooks with prospective, equity-aware evaluation. Parallel tracks should advance LLM alignment with KG guardrails, publish model cards and change protocols, and establish pathways for regulatory readiness and post-market monitoring. With shared standards, open resources, and rigorous clinical validation, the community can translate the KG framework from concept to practice delivering earlier detection, more precise and equitable treatment selection, and durable improvements in pediatric sleep health.

Author Contributions: J.D.G., J.G., P.G.C., P.R.C., R.J.H.-M., C.R.G., A.I.Q., S.G.B. and S.R.G. All authors of the manuscript meet all of the following conditions: (1) Substantial contributions to the conception or design of the work; or the acquisition, analysis, or interpretation of data for the work; AND (2) Drafting the work or reviewing it critically for important intellectual content; AND (3) Final approval of the version to be published; AND (4) Agreement to be accountable for all aspects of the work in ensuring that questions related to the accuracy or integrity of any part of the work are appropriately investigated and resolved. Conceptualization, J.D.G., J.G., P.G.C., P.R.C., C.R.G. and A.I.Q.; methodology, J.D.G., J.G., P.G.C. and P.R.C.; software, J.D.G., J.G. and P.G.C.; validation, J.D.G., J.G. and P.G.C.; formal analysis, J.D.G., J.G., P.G.C., R.J.H.-M. and S.R.G.; data and system curation, R.J.H.-M.; writing—original draft preparation, J.D.G., J.G., P.G.C. and S.G.B.; writing—critical review and editing, C.R.G., A.I.Q., R.J.H.-M., S.R.G. and P.R.C.; visualization, P.G.C. and C.R.G.; supervision, J.D.G., J.G. and P.G.C.; project administration, R.J.H.-M. Authorship was limited to those who contributed to the work reported. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: The investigations were conducted in accordance with the principles outlined in the Declaration of Helsinki (1975, revised in 2013). According to point 23 of this declaration, approval must be obtained from the local Institutional Review Board (IRB) or another relevant ethics committee prior to conducting the research, ensuring that the study adheres to both national and international guidelines. Ethics Committee Name: University of Alabama Human Research Protections Program Institutional Review Board. Approval Code: 24 01 7234. Approval Date: 21 March 2024.

Informed Consent Statement: Informed consent was waived for this study with fully deidentified data sources.

Data Availability Statement: No new data was reported in this manuscript.

Acknowledgments: During the preparation of this manuscript/study, the author(s) used Claude AI version 4 as a supplement to identify relevant references, ChatGPT 5 to assist creating the base outline for the document and checking for additional related topics and references primarily for the introduction and background components of the document, ChatGPT 5 was used to create the images used in the figures based on the conceptual design and text created by the authors, and Zotero to organize all references. The AI tools referenced (Claude AI and ChatGPT 5) were used solely to assist in generating a preliminary high-level outline and to suggest general background topics for consideration during manuscript development. These outputs were used as a starting point only and were substantially modified, expanded, and validated by the authors. No AI-generated text was incorporated directly into the final manuscript without critical revision, verification, and authorship oversight. All scientific content, interpretation, and conclusions presented in the manuscript are the original work of the authors. The authors confirm that all references were manually verified by the authors. The manuscript was written by the authors. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: P.R.C. receives book royalties from McGraw-Hill Publishers, Wolters Kluwer/Lippincott Williams & Wilkins Publishers, and CRC Press/Taylor & Francis. P.R.C. is a co-founder of Chronobionics, LLC; Neurotechion, LLC; and Biom Therapeutics, LLC. P.R.C. serves as Editor-in-Chief of Children (MDPI), the journal to which this manuscript is submitted; he had no role in the peer review or editorial decision for this article. J.D.G. receives book royalties from McGraw-Hill Publishers and Wolters Kluwer/Lippincott Williams & Wilkins Publishers. J.D.G. is a cofounder of Magic Medical Solutions, LLC, Evimero, LLC, and serves on the board of Pression, LLC. J.D.G receives funding through grants and contracts from the NSF, NIH, and ARPA-H (research unrelated to the present manuscript). J.G receives funding through grants and contracts from the NSF and NIH (research unrelated to the present manuscript). P.G.C. is a cofounder of Magic Medical Solutions, LLC and Evimero, LLC. C.R.G. is a co-founder of C&K Consulting. A.I.Q. is the co-founder of QureVasc, QureMed, and DyQure. All other authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AAP	American Academy of Pediatrics
ADHD	Attention-deficit/hyperactivity disorder
AHI	Apnea-hypopnea index
AI	Artificial intelligence
BPAP	Bilevel positive airway pressure
CDS	Clinical decision support
CNN	Convolutional neural network
CPAP	Continuous positive airway pressure
EHR	Electronic health record
ESS-CHAD	Epworth Sleepiness Scale for Children and Adolescents
FHIR	Fast Healthcare Interoperability Resources
GNN	Graph neural network
ICD-10-CM	International Classification of Diseases, Tenth Revision, Clinical Modification
KG	Knowledge graph
LLM	Large language model
LOINC	Logical Observation Identifiers Names and Codes
NLP	Natural language processing
OSA	Obstructive sleep apnea
OWL	Web Ontology Language
PSG	Polysomnography
PSQ	Pediatric Sleep Questionnaire
RAG	Retrieval-augmented generation
REM	Rapid eye movement
SHACL	Shapes Constraint Language
SNOMED CT	Systematized Nomenclature of Medicine—Clinical Terms
SPARQL	SPARQL Protocol and RDF Query Language
UCUM	Unified Code for Units of Measure

References

1. AASM | International Classification of Sleep Disorders. American Academy of Sleep Medicine—Association for Sleep Clinicians and Researchers. Available online: https://aasm.org/clinical-resources/international-classification-sleep-disorders/ (accessed on 19 February 2026).
2. Marcus, C.L.; Brooks, L.J.; Draper, K.A.; Gozal, D.; Halbower, A.C.; Jones, J.; Schechter, M.S.; Sheldon, S.H.; Spruyt, K.; Ward, S.D.; et al. Diagnosis and Management of Childhood Obstructive Sleep Apnea Syndrome. *Pediatrics* **2012**, *130*, 576–584. [CrossRef]
3. Aurora, R.N.; Zak, R.S.; Karippot, A.; Lamm, C.I.; Morgenthaler, T.I.; Auerbach, S.H.; Bista, S.R.; Casey, K.R.; Chowdhuri, S.; Kristo, D.A.; et al. Practice Parameters for the Respiratory Indications for Polysomnography in Children. *Sleep* **2011**, *34*, 379–388. [CrossRef] [PubMed] [PubMed Central]

4. Kaditis, A.G.; Alvarez, M.L.A.; Boudewyns, A.; Alexopoulos, E.I.; Ersu, R.; Joosten, K.; Larramona, H.; Miano, S.; Narang, I.; Trang, H.; et al. Obstructive sleep disordered breathing in 2- to 18-year-old children: Diagnosis and management. *Eur. Respir. J.* **2015**, *47*, 69–94. [CrossRef] [PubMed]
5. Soori, R.; Baikunje, N.; D'sa, I.; Bhushan, N.; Nagabhushana, B.; Hosmane, G.B. Pitfalls of AHI system of severity grading in obstructive sleep apnoea. *Sleep Sci.* **2022**, *15*, 285–288. [CrossRef] [PubMed] [PubMed Central]
6. Piotto, M.; Gambadauro, A.; Rocchi, A.; Lelii, M.; Madini, B.; Cerrato, L.; Chironi, F.; Belhaj, Y.; Patria, M.F. Pediatric Sleep Respiratory Disorders: A Narrative Review of Epidemiology and Risk Factors. *Children* **2023**, *10*, 955. [CrossRef] [PubMed] [PubMed Central]
7. Gouthro, K.; Kadam, S.J.; Sankari, A.; Slowik, J.M. Pediatric Obstructive Sleep Apnea. In *StatPearls*; StatPearls Publishing: Treasure Island, FL, USA, 2025. Available online: http://www.ncbi.nlm.nih.gov/books/NBK557610/ (accessed on 19 February 2026). [PubMed]
8. Spilsbury, J.C.; Storfer-Isser, A.; Rosen, C.L.; Redline, S. Remission and incidence of obstructive sleep apnea from middle childhood to late adolescence. *Sleep* **2015**, *38*, 23–29. [CrossRef] [PubMed] [PubMed Central]
9. Hasuneh, M.M.; Toubasi, A.A.; Khraisat, B.; Aldabbas, H.; Al-Iede, M. Risk Factors of Obstructive Sleep Apnea (OSA) in Pediatric Patients: A Systematic Review and Meta-analysis. *J. Pediatr. Health Care* **2024**, *38*, 717–726. [CrossRef] [PubMed]
10. Kang, K.T.; Weng, W.C.; Lee, P.L.; Hsu, W.C. Age- and gender-related characteristics in pediatric obstructive sleep apnea. *Pediatr. Pulmonol.* **2022**, *57*, 1520–1526. [CrossRef] [PubMed]
11. Perfect, M.M.; Archbold, K.; Goodwin, J.L.; Levine-Donnerstein, D.; Quan, S.F. Risk of Behavioral and Adaptive Functioning Difficulties in Youth with Previous and Current Sleep Disordered Breathing. *Sleep* **2013**, *36*, 517–525. [CrossRef] [PubMed] [PubMed Central]
12. O'Brien, L.M.; Mervis, C.B.; Holbrook, C.R.; Bruner, J.L.; Smith, N.H.; McNally, N.; McClimment, M.C.; Gozal, D. Neurobehavioral correlates of sleep-disordered breathing in children. *J. Sleep Res.* **2004**, *13*, 165–172. [CrossRef] [PubMed]
13. Chervin, R.D.; Hedger, K.; Dillon, J.E.; Pituch, K.J. Pediatric sleep questionnaire (PSQ): Validity and reliability of scales for sleep-disordered breathing, snoring, sleepiness, and behavioral problems. *Sleep Med.* **2000**, *1*, 21–32. [CrossRef] [PubMed]
14. Savini, S.; Ciorba, A.; Bianchini, C.; Stomeo, F.; Corazzi, V.; Vicini, C.; Pelucchi, S. Assessment of obstructive sleep apnoea (OSA) in children: An update. *Acta Otorhinolaryngol. Ital.* **2019**, *39*, 289–297. [CrossRef] [PubMed] [PubMed Central]
15. Unchiti, K.; Samerchua, A.; Pipanmekaporn, T.; Leurcharusmee, P.; Sonsuwan, N.; Phinyo, P.; Patumanond, J. Clinical predictors of moderate-to-severe pediatric obstructive sleep apnea. *Front. Pediatr.* **2024**, *12*, 1421467. [CrossRef]
16. Stierer, T.L. Big data: Mind your Ps and Qs—The importance of precision and quality in data validation. *J. Clin. Sleep Med.* **2020**, *16*, 161–162. [CrossRef] [PubMed] [PubMed Central]
17. Davenport, M.A.; Sirrianni, J.W.; Chisolm, D.J. Machine learning data sources in pediatric sleep research: Assessing racial/ethnic differences in electronic health record–based clinical notes prior to model training. *Front. Sleep* **2024**, *3*, 1271167. [CrossRef]
18. Boira, I.; Sancho-Chust, J.N.; Esteban, V.; Pastor, E.; Martínez-García, M.Á.; Chiner, E. Cross-Cultural Adaptation of a Health-Related Quality-of-Life Questionnaire for Children with Obstructive Sleep Disorders: Spanish Version of the OSD-6. *J. Clin. Med.* **2025**, *14*, 1709. [CrossRef]
19. Amann, J.; Vetter, D.; Blomberg, S.N.; Christensen, H.C.; Coffee, M.; Gerke, S.; Gilbert, T.K.; Hagendorff, T.; Holm, S.; Livne, M.; et al. To explain or not to explain?—Artificial intelligence explainability in clinical decision support systems. *PLoS Digit. Health* **2022**, *1*, e0000016. [CrossRef] [PubMed] [PubMed Central]
20. Gutiérrez-Tobal, G.C.; Álvarez, D.; Kheirandish-Gozal, L.; del Campo, F.; Gozal, D.; Hornero, R. Reliability of machine learning to diagnose pediatric obstructive sleep apnea: Systematic review and meta-analysis. *Pediatr. Pulmonol.* **2022**, *57*, 1931–1943. [CrossRef] [PubMed] [PubMed Central]
21. Al Khatib, H.S.; Neupane, S.; Kumar Manchukonda, H.; Golilarz, N.A.; Mittal, S.; Amirlatifi, A.; Rahimi, S. Patient-centric knowledge graphs: A survey of current methods, challenges, and applications. *Front. Artif. Intell.* **2024**, *7*, 1388479. [CrossRef] [PubMed] [PubMed Central]
22. Murali, L.; Gopakumar, G.; Viswanathan, D.M.; Nedungadi, P. Towards electronic health record-based medical knowledge graph construction, completion, and applications: A literature study. *J. Biomed. Inform.* **2023**, *143*, 104403. [CrossRef] [PubMed]
23. Das, S.; Hussey, P. HL7-FHIR-Based ContSys Formal Ontology for Enabling Continuity of Care Data Interoperability. *J. Pers. Med.* **2023**, *13*, 1024. [CrossRef]
24. Shang, Y.; Tian, Y.; Lyu, K.; Zhou, T.; Zhang, P.; Chen, J.; Li, J. Electronic Health Record–Oriented Knowledge Graph System for Collaborative Clinical Decision Support Using Multicenter Fragmented Medical Data: Design and Application Study. *J. Med. Internet Res.* **2024**, *26*, e54263. [CrossRef] [PubMed]
25. Alam, F.; Giglou, H.B.; Malik, K.M. Automated clinical knowledge graph generation framework for evidence based medicine. *Expert Syst. Appl.* **2023**, *233*, 120964. [CrossRef]
26. Mandel, J.C.; Kreda, D.A.; Mandl, K.D.; Kohane, I.S.; Ramoni, R.B. SMART on FHIR: A standards-based, interoperable apps platform for electronic health records. *J. Am. Med. Inform. Assoc.* **2016**, *23*, 899–908. [CrossRef] [PubMed] [PubMed Central]

27. BaHammam, A.S.; Pirzada, A.R.; Pandi-Perumal, S.R. Neurocognitive, mood changes, and sleepiness in patients with REM-predominant obstructive sleep apnea. *Sleep Breath.* **2023**, *27*, 57–66. [CrossRef] [PubMed]
28. Lee, S.; Yin, C.; Zhang, P. Stable clinical risk prediction against distribution shift in electronic health records. *Patterns* **2023**, *4*, 100828. [CrossRef] [PubMed] [PubMed Central]
29. Wu, C.R.; Tu, Y.K.; Chuang, L.P.; Gordon, C.; Chen, N.H.; Chen, P.Y.; Hasan, F.; Kurniasari, M.D.; Susanty, S.; Chiu, H.Y. Diagnostic meta-analysis of the Pediatric Sleep Questionnaire, OSA-18, and pulse oximetry in detecting pediatric obstructive sleep apnea syndrome. *Sleep Med. Rev.* **2020**, *54*, 101355. [CrossRef]
30. von Elm, E.; Altman, D.G.; Egger, M.; Pocock, S.J.; Gøtzsche, P.C.; Vandenbroucke, J.P. The Strengthening the Reporting of Observational Studies in Epidemiology (STROBE) statement: Guidelines for reporting observational studies. *J. Clin. Epidemiol.* **2008**, *61*, 344–349. [CrossRef] [PubMed]
31. Umano, G.R.; Rondinelli, G.; Luciano, M.; Pennarella, A.; Aiello, F.; Mangoni di Santo Stefano, G.S.R.C.; Di Sessa, A.; Marzuillo, P.; Papparella, A.; Miraglia del Giudice, E. Pediatric Sleep Questionnaire Predicts Moderate-to-Severe Obstructive Sleep Apnea in Children and Adolescents with Obesity. *Children* **2022**, *9*, 1303. [CrossRef] [PubMed] [PubMed Central]
32. Bhattacharjee, R.; Kheirandish-Gozal, L.; Spruyt, K.; Mitchell, R.B.; Promchiarak, J.; Simakajornboon, N.; Kaditis, A.G.; Splaingard, D.; Splaingard, M.; Brooks, L.J.; et al. Adenotonsillectomy outcomes in treatment of obstructive sleep apnea in children: A multicenter retrospective study. *Am. J. Respir. Crit. Care Med.* **2010**, *182*, 676–683. [CrossRef] [PubMed]
33. Ishman, S.L.; Maturo, S.; Schwartz, S.; McKenna, M.; Baldassari, C.M.; Bergeron, M.; Chernobilsky, B.; Ehsan, Z.; Gagnon, L.; Liu, Y.C.; et al. Expert Consensus Statement: Management of Pediatric Persistent Obstructive Sleep Apnea After Adenotonsillectomy. *Otolaryngol. Head Neck Surg.* **2023**, *168*, 115–130. [CrossRef] [PubMed] [PubMed Central]
34. Kang, K.T.; Chou, C.H.; Weng, W.C.; Lee, P.L.; Hsu, W.C. Associations between adenotonsillar hypertrophy, age, and obesity in children with obstructive sleep apnea. *PLoS ONE* **2013**, *8*, e78666. [CrossRef] [PubMed] [PubMed Central]
35. Cielo, C.M.; Taylor, J.A.; Vossough, A.; Radcliffe, J.; Thomas, A.; Bradford, R.; Lioy, J.; Tapia, I.E.; Assadsangabi, R.; Shults, J.; et al. Evolution of Obstructive Sleep Apnea in Infants with Cleft Palate and Micrognathia. *J. Clin. Sleep Med.* **2016**, *12*, 979–987. [CrossRef] [PubMed] [PubMed Central]
36. Schwartz, A.R.; Patil, S.P.; Laffan, A.M.; Polotsky, V.; Schneider, H.; Smith, P.L. Obesity and Obstructive Sleep Apnea. *Proc. Am. Thorac. Soc.* **2008**, *5*, 185–192. [CrossRef] [PubMed] [PubMed Central]
37. Dempsey, J.A.; Veasey, S.C.; Morgan, B.J.; O'Donnell, C.P. Pathophysiology of sleep apnea. *Physiol. Rev.* **2010**, *90*, 47–112. [CrossRef] [PubMed] [PubMed Central]
38. Katz, E.S.; D'Ambrosio, C.M. Pathophysiology of pediatric obstructive sleep apnea. *Proc. Am. Thorac. Soc.* **2008**, *5*, 253–262. [CrossRef] [PubMed] [PubMed Central]
39. Chirakalwasan, N.; Ruxrungtham, K. The linkage of allergic rhinitis and obstructive sleep apnea. *Asian Pac. J. Allergy Immunol.* **2014**, *32*, 276–286. [PubMed]
40. Wang, J.; Zhao, Y.; Yang, W.; Shen, T.; Xue, P.; Yan, X.; Chen, D.; Qiao, Y.; Chen, M.; Ren, R.; et al. Correlations between obstructive sleep apnea and adenotonsillar hypertrophy in children of different weight status. *Sci. Rep.* **2019**, *9*, 11455. [CrossRef]
41. Muzumdar, H.; Arens, R. Diagnostic Issues in Pediatric Obstructive Sleep Apnea. *Proc. Am. Thorac. Soc.* **2008**, *5*, 263–273. [CrossRef] [PubMed] [PubMed Central]
42. Tan, H.L.; Gozal, D.; Kheirandish-Gozal, L. Obstructive sleep apnea in children: A critical update. *Nat. Sci. Sleep* **2013**, *5*, 109–123. [CrossRef] [PubMed]
43. Chervin, R.D.; Archbold, K.H.; Dillon, J.E.; Panahi, P.; Pituch, K.J.; Dahl, R.E.; Guilleminault, C. Inattention, hyperactivity, and symptoms of sleep-disordered breathing. *Pediatrics* **2002**, *109*, 449–456. [CrossRef] [PubMed]
44. Lumeng, J.C.; Chervin, R.D. Epidemiology of Pediatric Obstructive Sleep Apnea. *Proc. Am. Thorac. Soc.* **2008**, *5*, 242–252. [CrossRef] [PubMed] [PubMed Central]
45. Nosetti, L.; Gozal, D. Exploring the bidirectional relationship between asthma and obstructive sleep apnea in Brazilian pediatric patients: One more piece to the Puzzle. *J. Pediatr.* **2023**, *99*, 423–424. [CrossRef] [PubMed] [PubMed Central]
46. Garza, N.; Witmans, M.; Salud, M.; Lagera, P.G.D.; Co, V.A.; Tablizo, M.A. The Association between Asthma and OSA in Children. *Children* **2022**, *9*, 1430. [CrossRef]
47. Smith, D.F.; Amin, R.S. OSA and Cardiovascular Risk in Pediatrics. *Chest* **2019**, *156*, 402–413. [CrossRef] [PubMed] [PubMed Central]
48. Costa, D.J.; Mitchell, R. Adenotonsillectomy for obstructive sleep apnea in obese children: A meta-analysis. *Otolaryngol. Head Neck Surg.* **2009**, *140*, 455–460. [CrossRef] [PubMed]
49. Urbano, G.L.; Tablizo, B.J.; Moufarrej, Y.; Tablizo, M.A.; Chen, M.L.; Witmans, M. The Link between Pediatric Obstructive Sleep Apnea (OSA) and Attention Deficit Hyperactivity Disorder (ADHD). *Children* **2021**, *8*, 824. [CrossRef] [PubMed] [PubMed Central]

50. Kim, J.; Bhattacharjee, R.; Dayyat, E.; Snow, A.B.; Kheirandish-Gozal, L.; Goldman, J.L.; Li, R.C.; Serpero, L.D.; Clair, H.B.; Gozal, D. Increased cellular proliferation and inflammatory cytokines in tonsils derived from children with obstructive sleep apnea. *Pediatr. Res.* **2009**, *66*, 423–428. [CrossRef] [PubMed] [PubMed Central]
51. Burton, M.J.; Glasziou, P.P.; Chong, L.Y.; Venekamp, R.P. Tonsillectomy or adenotonsillectomy versus non-surgical treatment for chronic/recurrent acute tonsillitis. *Cochrane Database Syst. Rev.* **2014**, *2014*, CD001802. [CrossRef] [PubMed] [PubMed Central]
52. Beck, S.E.; Marcus, C.L. Pediatric Polysomnography. *Sleep Med. Clin.* **2009**, *4*, 393–406. [CrossRef] [PubMed] [PubMed Central]
53. Stowe, R.C.; Afolabi-Brown, O. Pediatric polysomnography—A review of indications, technical aspects, and interpretation. *Paediatr. Respir. Rev.* **2020**, *34*, 9–17. [CrossRef] [PubMed]
54. Janssen, K.C.; Phillipson, S.; O'Connor, J.; Johns, M.W. Validation of the Epworth Sleepiness Scale for Children and Adolescents using Rasch analysis. *Sleep Med.* **2017**, *33*, 30–35. [CrossRef] [PubMed]
55. King, M.S.; Xanthopoulos, M.S.; Marcus, C.L. Improving Positive Airway Pressure Adherence in Children. *Sleep Med. Clin.* **2014**, *9*, 219–234. [CrossRef] [PubMed] [PubMed Central]
56. Bahammam, S.A. Rapid Maxillary Expansion for Obstructive Sleep Apnea among children—Systematic Review and Meta-analysis. *Sleep Sci.* **2020**, *13*, 70–77. [PubMed] [PubMed Central]
57. Liming, B.J.; Ryan, M.; Mack, D.; Ahmad, I.; Camacho, M. Montelukast and Nasal Corticosteroids to Treat Pediatric Obstructive Sleep Apnea: A Systematic Review and Meta-analysis. *Otolaryngol. Head Neck Surg.* **2019**, *160*, 594–602. [CrossRef] [PubMed]
58. Gupta, S.; Sharma, R. Pediatric Obstructive Sleep Apnea: Diagnostic Challenges and Management Strategies. *Cureus* **2024**, *16*, e75347. [CrossRef] [PubMed] [PubMed Central]
59. Phan, H.; Mertins, A.; Baumert, M. Pediatric Automatic Sleep Staging: A Comparative Study of State-of-the-Art Deep Learning Methods. *IEEE Trans. Biomed. Eng.* **2022**, *69*, 3612–3622. [CrossRef] [PubMed]
60. Somaskandhan, P.; Leppänen, T.; Terrill, P.I.; Sigurdardottir, S.; Arnardottir, E.S.; Ólafsdóttir, K.A.; Serwatko, M.; Sigurðardóttir, S.Þ.; Clausen, M.; Töyräs, J.; et al. Deep learning-based algorithm accurately classifies sleep stages in preadolescent children with sleep-disordered breathing symptoms and age-matched controls. *Front. Neurol.* **2023**, *14*, 1162998. [CrossRef] [PubMed]
61. Pandey, S.R.; Saeed, A.; Lee, H. PedSleepMAE: Generative Model for Multimodal Pediatric Sleep Signals. In Proceedings of the 2024 IEEE EMBS International Conference on Biomedical and Health Informatics (BHI), Houston, TX, USA, 10–13 November 2024; pp. 1–8. [CrossRef]
62. Qin, H.; Zhang, L.; Li, X.; Xu, Z.; Zhang, J.; Wang, S.; Zheng, L.; Ji, T.; Mei, L.; Kong, Y.; et al. Pediatric obstructive sleep apnea diagnosis: Leveraging machine learning with linear discriminant analysis. *Front. Pediatr.* **2024**, *12*, 1328209. [CrossRef] [PubMed] [PubMed Central]
63. Liu, X.; Pamula, Y.; Immanuel, S.; Kennedy, D.; Martin, J.; Baumert, M. Utilisation of machine learning to predict surgical candidates for the treatment of childhood upper airway obstruction. *Sleep Breath.* **2022**, *26*, 649–661. [CrossRef]
64. Han, H.; Oh, J. Application of various machine learning techniques to predict obstructive sleep apnea syndrome severity. *Sci. Rep.* **2023**, *13*, 6379. [CrossRef]
65. Wang, Z.; Bao, X.; Zhao, C.; Zhang, J.; Ai, S.; Li, Y. UT-OSANet: A Multimodal Deep Learning model for Evaluating and Classifying Obstructive Sleep Apnea. *arXiv* **2025**, arXiv:2511.16169. [CrossRef]
66. Bazoukis, G.; Bollepalli, S.C.; Chung, C.T.; Li, X.; Tse, G.; Bartley, B.L.; Batool-Anwar, S.; Quan, S.F.; Armoundas, A.A. Application of artificial intelligence in the diagnosis of sleep apnea. *J. Clin. Sleep Med.* **2023**, *19*, 1337–1363. [CrossRef] [PubMed] [PubMed Central]
67. Liu, X.; Glocker, B.; McCradden, M.M.; Ghassemi, M.; Denniston, A.K.; Oakden-Rayner, L. The medical algorithmic audit. *Lancet Digit. Health* **2022**, *4*, e384–e397. [CrossRef] [PubMed]
68. Adamusiak, T.; Shimoyama, N.; Shimoyama, M. Next Generation Phenotyping Using the Unified Medical Language System. *JMIR Med. Inform.* **2014**, *2*, e5. [CrossRef] [PubMed] [PubMed Central]
69. Aldughayfiq, B.; Ashfaq, F.; Jhanjhi, N.Z.; Humayun, M. Capturing Semantic Relationships in Electronic Health Records Using Knowledge Graphs: An Implementation Using MIMIC III Dataset and GraphDB. *Healthcare* **2023**, *11*, 1762. [CrossRef] [PubMed] [PubMed Central]
70. Kahn, M.G.; Bailey, L.C.; Forrest, C.B.; Padula, M.A.; Hirschfeld, S. Building a Common Pediatric Research Terminology for Accelerating Child Health Research. *Pediatrics* **2014**, *133*, 516–525. [CrossRef] [PubMed] [PubMed Central]
71. Hu, Y.; Xu, C.; Lin, B.; Yang, W.; Tang, Y.Y. Medical multimodal large language models: A systematic review. *Intell. Oncol.* **2025**, *1*, 308–325. [CrossRef]
72. Jhee, J.H.; Megina, A.; Beaufils, P.C.D.; Karakachoff, M.; Redon, R.; Gaignard, A.; Coulet, A. Predicting clinical outcomes from patient care pathways represented with temporal knowledge graphs. *arXiv* **2025**, arXiv:2502.21138. [CrossRef]
73. Zhu, Y.; Ren, C.; Xie, S.; Liu, S.; Ji, H.; Wang, Z.; Sun, T.; He, L.; Li, Z.; Zhu, X.; et al. REALM: RAG-Driven Enhancement of Multimodal Electronic Health Records Analysis via Large Language Models. *arXiv* **2024**, arXiv:2402.07016. [CrossRef]

74. Boxwala, A.A.; Rocha, B.H.; Maviglia, S.; Kashyap, V.; Meltzer, S.; Kim, J.; Tsurikova, R.; Wright, A.; Paterno, M.D.; Fairbanks, A.; et al. A multi-layered framework for disseminating knowledge for computer-based decision support. *J. Am. Med. Inform. Assoc.* **2011**, *18*, i132–i139. [CrossRef] [PubMed] [PubMed Central]
75. Section on Pediatric Pulmonology; Subcommittee on Obstructive Sleep Apnea Syndrome. Clinical Practice Guideline: Diagnosis and Management of Childhood Obstructive Sleep Apnea Syndrome. *Pediatrics* **2002**, *109*, 704–712. [CrossRef]
76. Kim, K.A.; Kim, S.J.; Yoon, A. Craniofacial anatomical determinants of pediatric sleep-disordered breathing: A comprehensive review. *J. Prosthodont.* **2025**, *34*, 26–34. [CrossRef] [PubMed]
77. He, S.; Li, Y.; Zhang, C.; Li, Z.; Ren, Y.; Li, T.; Wang, J. Deep learning technique to detect craniofacial anatomical abnormalities concentrated on middle and anterior of face in patients with sleep apnea. *Sleep Med.* **2023**, *112*, 12–20. [CrossRef]
78. Marcus, C.L.; Moore, R.H.; Rosen, C.L.; Giordani, B.; Garetz, S.L.; Taylor, H.G.; Mitchell, R.B.; Amin, R.; Katz, E.S.; Arens, R.; et al. A randomized trial of adenotonsillectomy for childhood sleep apnea. *N. Engl. J. Med.* **2013**, *368*, 2366–2376. [CrossRef] [PubMed] [PubMed Central]
79. Wang, Y.G.; Menno, D.; Chen, A.; Steininger, T.L.; Morris, S.; Black, J.; Profant, J.; Johns, M.W. Validation of the Epworth Sleepiness Scale for Children and Adolescents (ESS-CHAD) questionnaire in pediatric patients with narcolepsy with cataplexy aged 7–16 years. *Sleep Med.* **2022**, *89*, 78–84. [CrossRef] [PubMed]
80. Bodenreider, O.; Cornet, R.; Vreeman, D.J. Recent Developments in Clinical Terminologies—SNOMED CT, LOINC, and RxNorm. *Yearb. Med. Inform.* **2018**, *27*, 129–139. [CrossRef] [PubMed] [PubMed Central]
81. McMurry, A.J.; Gottlieb, D.I.; Miller, T.A.; Jones, J.R.; Atreja, A.; Crago, J.; Desai, P.M.; Dixon, B.E.; Garber, M.; Ignatov, V.; et al. Cumulus: A federated EHR-based learning system powered by FHIR and AI. *medRxiv* **2024**, 2024.02.02.24301940. [CrossRef] [PubMed] [PubMed Central]
82. Kang, M.; Mo, F.; Witmans, M.; Santiago, V.; Tablizo, M.A. Trends in Diagnosing Obstructive Sleep Apnea in Pediatrics. *Children* **2022**, *9*, 306. [CrossRef] [PubMed] [PubMed Central]
83. Teplitzky, T.B.; Zauher, A.; Isaiah, A. Evaluation and diagnosis of pediatric obstructive sleep apnea—An update. *Front. Sleep* **2023**, *2*, 1127784. [CrossRef]
84. Patel, V.P.; Patroneva, A.; Glaze, D.G.; Davis, M.S.K.; Merikle, E.; Revana, A. Establishing the content validity of the Epworth Sleepiness Scale for Children and Adolescents in Prader-Willi syndrome. *J. Clin. Sleep Med.* **2022**, *18*, 485–496. [CrossRef] [PubMed]
85. Chuang, H.-H.; Huang, C.-G.; Chuang, L.-P.; Huang, Y.-S.; Chen, N.-H.; Li, H.-Y.; Fang, T.-J.; Hsu, J.-F.; Lai, H.-C.; Chen, J.-Y.; et al. Relationships Among and Predictive Values of Obesity, Inflammation Markers, and Disease Severity in Pediatric Patients with Obstructive Sleep Apnea Before and After Adenotonsillectomy. *J. Clin. Med.* **2020**, *9*, 579. [CrossRef] [PubMed] [PubMed Central]
86. U.S. National Library of Medicine. SNOMED CT [Product Program Project Descriptions]. Available online: https://www.nlm.nih.gov/healthit/snomedct/index.html (accessed on 25 February 2026).
87. Cremonini, F.; Zucchini, L.; Pellitteri, F.; Palone, M.; Lombardo, L. Obstructive Sleep Apnea in Developmental Age: 22-Item Pediatric Sleep Questionnaire for an Observational Descriptive Investigation. *Children* **2023**, *10*, 1265. [CrossRef] [PubMed]
88. Domany, K.A.; Dana, E.; Tauman, R.; Gut, G.; Greenfeld, M.; Yakir, B.E.; Sivan, Y. Adenoidectomy for Obstructive Sleep Apnea in Children. *J. Clin. Sleep Med.* **2016**, *12*, 1285–1291. [CrossRef]
89. Wołk, K. Evaluating Retrieval-Augmented Generation Variants for Clinical Decision Support: Hallucination Mitigation and Secure On-Premises Deployment. *Electronics* **2025**, *14*, 4227. [CrossRef]
90. Zhou, S.; Huang, M.; Cole, J.; Britton, C.; Yin, M.; Wolber, J.; Li, K. Assessing Automated Fact-Checking for Medical LLM Responses with Knowledge Graphs. *Proc. AAAI Conf. Artif. Intell.* **2026**, *40*, 39646–39656. [CrossRef]
91. Hao, Y.; Cheng, C.; Li, J.; Li, H.; Di, X.; Zeng, X.; Jin, S.; Han, X.; Liu, C.; Wang, Q.; et al. Multimodal Integration in Health Care: Development with Applications in Disease Management. *J. Med. Internet Res.* **2025**, *27*, e76557. [CrossRef]
92. SHACL and OWL Compared. Available online: https://spinrdf.org/shacl-and-owl.html (accessed on 25 February 2026).
93. Das, U.; Atmakuri, K.B.; Ho, D.; Lee, C.; Lee, Y. Clinical Knowledge Graph Construction and Evaluation with Multi-LLMs via Retrieval-Augmented Generation. *arXiv* **2026**, arXiv:2601.01844. [CrossRef]
94. Cui, H.; Lu, J.; Xu, R.; Wang, S.; Ma, W.; Yu, Y.; Yu, S.; Kan, S.; Ling, S.; Zhao, L.; et al. A review on knowledge graphs for healthcare: Resources, applications, and promises. *J. Biomed. Inform.* **2025**, *169*, 104861. [CrossRef] [PubMed]
95. Min, J.; Zhang, X.; Griffis, H.M.; Cielo, C.M.; Tapia, I.E.; Williamson, A.A. Sociodemographic Disparities and Healthcare Utilization in Pediatric Obstructive Sleep Apnea Management. *Sleep Med.* **2023**, *109*, 211–218. [CrossRef] [PubMed] [PubMed Central]
96. Johnson, D.A.; Ohanele, C.; Alcántara, C.; Jackson, C.L. The Need for Social and Environmental Determinants of Health Research to Understand and Intervene upon Racial/Ethnic Disparities in Obstructive Sleep Apnea. *Clin. Chest Med.* **2022**, *43*, 199–216. [CrossRef] [PubMed] [PubMed Central]
97. Vickers, A.J.; Elkin, E.B. Decision curve analysis: A novel method for evaluating prediction models. *Med. Decis. Mak.* **2006**, *26*, 565–574. [CrossRef] [PubMed] [PubMed Central]

98. McQuillan, M.E.; Jones, I.C.; Abu Mayyaleh, H.F.; Khan, S.; Honaker, S.M. Health disparities in the detection and prevalence of pediatric obstructive sleep apnea. *Front. Sleep* **2023**, *2*, 1225808. [CrossRef] [PubMed]
99. Keamy, D.G.; Chhabra, K.R.; Hartnick, C.J. Predictors of complications following adenotonsillectomy in children with severe obstructive sleep apnea. *Int. J. Pediatr. Otorhinolaryngol.* **2015**, *79*, 1838–1841. [CrossRef] [PubMed]
100. Ehsan, Z.; Ishman, S.L.; Soghier, I.; Almeida, F.R.; Boudewyns, A.; Camacho, M.; Carno, M.A.; Coppelson, K.; Ersu, R.H.; Ho, A.T.N.; et al. Management of Persistent, Post-adenotonsillectomy Obstructive Sleep Apnea in Children: An Official American Thoracic Society Clinical Practice Guideline. *Am. J. Respir. Crit. Care Med.* **2024**, *209*, 248–261. [CrossRef] [PubMed] [PubMed Central]
101. Noll, R.; Schaaf, J.; Storf, H. The Use of Computer-Assisted Case-Based Reasoning to Support Clinical Decision-Making—A Scoping Review. In *Case-Based Reasoning Research and Development*; Springer: Cham, Switzerland, 2022; pp. 395–409. [CrossRef]
102. Zhao, X.; Liu, S.; Yang, S.Y.; Miao, C. MedRAG: Enhancing Retrieval-augmented Generation with Knowledge Graph-Elicited Reasoning for Healthcare Copilot. In *WWW' 25: Proceedings of the ACM on Web Conference 2025*; Association for Computing Machinery: New York, NY, USA, 2025; pp. 4442–4457. [CrossRef]
103. Hossny, E.; Adachi, Y.; Anastasiou, E.; Badellino, H.; Custovic, A.; El-Owaidy, R.; El-Sayed, Z.A.; Filipovic, I.; Gomez, R.M.; Kalayci, O.; et al. Pediatric asthma comorbidities: Global impact and unmet needs. *World Allergy Organ. J.* **2024**, *17*, 100909. [CrossRef] [PubMed]
104. Wu, Y.; Zheng, L.; Cui, G.; Xu, Z.; Ni, X. Subtypes of obstructive sleep apnea in children and related factors. *J. Clin. Sleep Med.* **2022**, *18*, 2397–2404. [CrossRef]
105. D'Elia, C.; Gozal, D.; Bruni, O.; Goudouris, E.; Meira e Cruz, M. Allergic rhinitis and sleep disorders in children—Coexistence and reciprocal interactions. *J. Pediatr.* **2021**, *98*, 444–454. [CrossRef] [PubMed] [PubMed Central]
106. Visweswaran, S.; Zhang, L.Y.; Bui, K.; Sadhu, E.M.; Samayamuthu, M.J.; Morris, M.M. Sharing and Reusing Computable Phenotype Definitions. *medRxiv* **2023**, 2023.09.17.23295681. [CrossRef]
107. Richesson, R.L.; Smerek, M.M.; Blake Cameron, C. A Framework to Support the Sharing and Reuse of Computable Phenotype Definitions Across Health Care Delivery and Clinical Research Applications. *eGEMS* **2016**, *4*, 1232. [CrossRef] [PubMed] [PubMed Central]
108. Gutiérrez-Tobal, G.C.; Gomez-Pilar, J.; Ferreira-Santos, D.; Pereira-Rodrigues, P.; Álvarez, D.; Campo, F.; Gozal, D.; Hornero, R. A subject-based association network defines new pediatric sleep apnea phenotypes with different odds of recovery after treatment. *Comput. Methods Programs Biomed.* **2026**, *275*, 109209. [CrossRef] [PubMed]
109. Ma, E.Y.; Kim, J.W.; Lee, Y.; Cho, S.W.; Kim, H.; Kim, J.K. Combined unsupervised-supervised machine learning for phenotyping complex diseases with its application to obstructive sleep apnea. *Sci. Rep.* **2021**, *11*, 4457. [CrossRef] [PubMed] [PubMed Central]
110. Zinchuk, A.V.; Gentry, M.J.; Concato, J.; Yaggi, H.K. Phenotypes in obstructive sleep apnea: A definition, examples and evolution of approaches. *Sleep Med. Rev.* **2017**, *35*, 113–123. [CrossRef] [PubMed] [PubMed Central]
111. Wang, Z.; Shi, D.; Zhao, J.; Diao, X.; Tang, X.; Qin, Y. Automated Construction of Medical Indicator Knowledge Graphs Using Retrieval Augmented Large Language Models. *arXiv* **2025**, arXiv:2511.13526. [CrossRef]
112. He, J.; Zhang, B.; Rouhizadeh, H.; Chen, Y.; Yang, R.; Lu, J.; Chen, X.; Liu, N.; Teodoro, D. Retrieval-Augmented Generation in Biomedicine: A Survey of Technologies, Datasets, and Clinical Applications. *arXiv* **2025**, arXiv:2505.01146. [CrossRef]
113. Schäfer, H.; Idrissi-Yaghir, A.; Arzideh, K.; Damm, H.; Pakull, T.M.G.; Schmidt, C.S.; Bahn, M.; Lodde, G.; Livingstone, E.; Schadendorf, D.; et al. BioKGrapher: Initial evaluation of automated knowledge graph construction from biomedical literature. *Comput. Struct. Biotechnol. J.* **2024**, *24*, 639–660. [CrossRef] [PubMed]
114. Qin, G.; Narsinh, K.; Wei, Q.; Roach, J.C.; Joshi, A.; Goetz, S.L.; Moxon, S.T.; Brush, M.H.; Xu, X.; Yao, Y.; et al. Generating Biomedical Knowledge Graphs from Knowledge Bases, Registries, and Multiomic Data. *bioRxiv* **2024**, 2024.11.14.623648. [CrossRef] [PubMed] [PubMed Central]
115. Oss Boll, H.; Amirahmadi, A.; Ghazani, M.M.; de Morais, W.O.; de Freitas, E.P.; Soliman, A.; Etminani, F.; Byttner, S.; Recamonde-Mendoza, M. Graph neural networks for clinical risk prediction based on electronic health records: A survey. *J. Biomed. Inform.* **2024**, *151*, 104616. [CrossRef] [PubMed]
116. Sawunyavisuth, B.; Ngamjarus, C.; Sawanyawisuth, K. Adherence to Continuous Positive Airway Pressure Therapy in Pediatric Patients with Obstructive Sleep Apnea: A Meta-Analysis. *Ther. Clin. Risk Manag.* **2023**, *19*, 143–162. [CrossRef] [PubMed] [PubMed Central]
117. Huang, Y.; Li, W.; Macheret, F.; Gabriel, R.A.; Ohno-Machado, L. A tutorial on calibration measurements and calibration models for clinical prediction models. *J. Am. Med. Inform. Assoc.* **2020**, *27*, 621–633. [CrossRef] [PubMed]
118. Ying, R.; Bourgeois, D.; You, J.; Zitnik, M.; Leskovec, J. GNNExplainer: Generating Explanations for Graph Neural Networks. *Adv. Neural Inf. Process Syst.* **2019**, *32*, 9240–9251. [PubMed] [PubMed Central]
119. Luo, D.; Cheng, W.; Xu, D.; Yu, W.; Zong, B.; Chen, H.; Zhang, X. Parameterized explainer for graph neural network. In *NIPS' 20: Proceedings of the 34th International Conference on Neural Information Processing Systems*; Curran Associates Inc.: Red Hook, NY, USA, 2020; pp. 19620–19631. Available online: https://dl.acm.org/doi/10.5555/3495724.3497370 (accessed on 26 February 2026).

120. As'ad, M.; Faran, N.; Joharji, H. AI-Supported Shared Decision-Making (AI-SDM): Conceptual Framework. *JMIR AI* **2025**, *4*, e75866. [CrossRef] [PubMed] [PubMed Central]
121. Gerdes, A. The role of explainability in AI-supported medical decision-making. *Discov. Artif. Intell.* **2024**, *4*, 29. [CrossRef]
122. Lu, C.; Angelopoulos, A.N.; Pomerantz, S. Improving Trustworthiness of AI Disease Severity Rating in Medical Imaging with Ordinal Conformal Prediction Sets. In *Medical Image Computing and Computer Assisted Intervention—MICCAI*; Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S., Eds.; Springer Nature: Cham, Switzerland, 2022; pp. 545–554. [CrossRef]
123. Rotmensch, M.; Halpern, Y.; Tlimat, A.; Horng, S.; Sontag, D. Learning a Health Knowledge Graph from Electronic Medical Records. *Sci. Rep.* **2017**, *7*, 5994. [CrossRef] [PubMed]
124. Char, D.S.; Shah, N.H.; Magnus, D. Implementing Machine Learning in Health Care—Addressing Ethical Challenges. *N. Engl. J. Med.* **2018**, *378*, 981–983. [CrossRef]
125. Rosenberg, R.S.; Van Hout, S. The American Academy of Sleep Medicine Inter-scorer Reliability Program: Sleep Stage Scoring. *J. Clin. Sleep Med.* **2013**, *9*, 81–87. [CrossRef] [PubMed]
126. Grigg-Damberger, M.M. The AASM Scoring Manual Four Years Later. *J. Clin. Sleep Med.* **2012**, *8*, 323–332. [CrossRef]
127. Finlayson, S.G.; Subbaswamy, A.; Singh, K.; Bowers, J.; Kupke, A.; Zittrain, J.; Kohane, I.S.; Saria, S. The Clinician and Dataset Shift in Artificial Intelligence. *N. Engl. J. Med.* **2021**, *385*, 283–286. [CrossRef] [PubMed] [PubMed Central]
128. Declerck, J.; Kılıç, Ö.D.; Erol, E.E.; Mehryar, S.; Kalra, D.; Ide, Z.; Celebi, R. Assessing Data Quality in Heterogeneous Health Care Integration: Simulation Study of the AIDAVA Framework. *JMIR Med. Inform.* **2025**, *13*, e75275. [CrossRef]
129. Kahn, M.G.; Callahan, T.J.; Barnard, J.; Bauck, A.E.; Brown, J.; Davidson, B.N.; Estiri, H.; Goerg, C.; Holve, E.; Johnson, S.G.; et al. A Harmonized Data Quality Assessment Terminology and Framework for the Secondary Use of Electronic Health Record Data. *eGEMS* **2016**, *4*, 1244. [CrossRef] [PubMed] [PubMed Central]
130. Ren, Z.; Wang, S.; Zhang, Y. Weakly supervised machine learning. *CAAI Trans. Intell. Technol.* **2023**, *8*, 549–580. [CrossRef]
131. Hansson, M.G.; Lochmüller, H.; Riess, O.; Schaefer, F.; Orth, M.; Rubinstein, Y.; Molster, C.; Dawkins, H.; Taruscio, D.; Posada, M.; et al. The risk of re-identification versus the need to identify individuals in rare disease research. *Eur. J. Hum. Genet.* **2016**, *24*, 1553–1558. [CrossRef] [PubMed] [PubMed Central]
132. Mawji, A.; Longstaff, H.; Trawin, J.; Dunsmuir, D.; Komugisha, C.; Novakowski, S.K.; Wiens, M.O.; Akech, S.; Tagoola, A.; Kissoon, N.; et al. A proposed de-identification framework for a cohort of children presenting at a health facility in Uganda. *PLoS Digit. Health* **2022**, *1*, e0000027. [CrossRef] [PubMed] [PubMed Central]
133. Ficek, J.; Wang, W.; Chen, H.; Dagne, G.; Daley, E. Differential privacy in health research: A scoping review. *J. Am. Med. Inform. Assoc.* **2021**, *28*, 2269–2276. [CrossRef] [PubMed] [PubMed Central]
134. Sheller, M.J.; Edwards, B.; Reina, G.A.; Martin, J.; Pati, S.; Kotrotsou, A.; Milchenko, M.; Xu, W.; Marcus, D.; Colen, R.R.; et al. Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data. *Sci. Rep.* **2020**, *10*, 12598. [CrossRef] [PubMed] [PubMed Central]
135. Pati, S.; Kumar, S.; Varma, A.; Edwards, B.; Lu, C.; Qu, L.; Wang, J.J.; Lakshminarayanan, A.; Wang, S.; Sheller, M.J.; et al. Privacy preservation for federated learning in health care. *Patterns* **2024**, *5*, 100974. [CrossRef] [PubMed] [PubMed Central]
136. Noy, N.F.; Shah, N.H.; Whetzel, P.L.; Dai, B.; Dorf, M.; Griffith, N.; Jonquet, C.; Rubin, D.L.; Storey, M.A.; Chute, C.G.; et al. BioPortal: Ontologies and integrated data resources at the click of a mouse. *Nucleic Acids Res.* **2009**, *37*, W170–W173. [CrossRef] [PubMed] [PubMed Central]
137. Cornet, R.; de Keizer, N. Forty years of SNOMED: A literature review. *BMC Med. Inform. Decis. Mak.* **2008**, *8*, S2. [CrossRef] [PubMed] [PubMed Central]
138. Morgan, K.L.; Kukhareva, P.V.; Warner, P.B.; Wilkof, J.; Snyder, M.; Horton, D.; Madsen, T.; Habboushe, J.; Kawamoto, K. Using CDS Hooks to increase SMART on FHIR app utilization: A cluster-randomized trial. *J. Am. Med. Inform. Assoc.* **2022**, *29*, 1461–1470. [CrossRef] [PubMed] [PubMed Central]

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.