
AGENTIC AI-ENABLED DISCOVERY ACROSS LARGE-SCALE SLEEP PHYSIOLOGY

Rahul Thapa^{1,2}, Umaer Hanif^{3,4,*}, Robin Guillard^{5,*}, Andreas Brink-Kjaer^{3,*}, Adrien Specht^{5,6,*}, Matteo Saibene^{3,*}, Magnus Ruud Kjaer^{3,4,5}, Harrison G. Zhang¹, Federico Bianchi⁷, Elisabeth Roxane M. Heremans⁵, Eric C. Landsness^{8,9}, Emmanuel Mignot^{5,+}, and James Zou^{1,2,+}

¹Department of Biomedical Data Science, Stanford University, Stanford, CA, USA

²Department of Computer Science, Stanford University, Stanford, CA, USA

³Department of Health Technology, Technical University of Denmark, Kongens Lyngby, Denmark

⁴Danish Center for Sleep Medicine, Department of Clinical Neurophysiology, Glostrup, Denmark

⁵Department of Psychiatry and Behavioral Sciences, Stanford University, Stanford, CA, USA

⁶Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA, USA

⁷Together AI, USA

⁸Department of Neurology, Washington University School of Medicine in St. Louis, MO, USA

⁹Consortium for Biomedical Research and AI in Neurodegeneration (C-BRAIN), USA

*Each led one case study and contributed equally. +Jointly supervised this work.

Correspondence: rthapa84@stanford.edu, mignot@stanford.edu, jamesz@stanford.edu

ABSTRACT

Sleep occupies roughly one-third of human life and plays a central role in human health, yet many aspects of its physiology remain poorly understood. Large polysomnography (PSG) datasets offer new opportunities to study disease risk, clinical phenotypes, and the biological organization of sleep, but extracting insight from these complex, multimodal recordings requires substantial expert effort and remains difficult for general-purpose AI systems. We therefore developed AI SLEEP Co-SCIENTIST, an expert-guided research environment in which human scientists direct specialist agents for hypothesis development, signal preprocessing, and statistical analysis, while reviewing intermediate outputs. Each reported result is linked to the executable code that produced it, providing an auditable record of the research process. Across four clinical and epidemiological cohorts comprising approximately 124,000 PSG recordings and more than 50 TB of raw signals, we conducted five case studies spanning three directions: how sleep physiology relates to future disease, how it distinguishes clinical phenotypes, and how sleep is organized and regulated. Diminished network-level physiological coupling during sleep was associated with incident Parkinson’s disease (HR 1.48, 95% CI 1.31–1.67) and Alzheimer’s disease (HR 1.38, 95% CI 1.25–1.53). A physiologically structured late-fusion sleep-age model outperformed an unconstrained early-fusion approach (mean absolute error 7.06 vs. 7.33 years in the Stanford Sleep Clinic cohort and 8.83 vs. 9.52 years in the Human Sleep Project), and its age residual was associated with incident disease across multiple organ systems. Arousal dynamics characterized comorbid insomnia and sleep apnoea as an intermediate phenotype skewed towards obstructive sleep apnoea, but distinguished from it by prolonged post-arousal wakefulness and more irregular arousal organization. Analyses of sleep regulation showed that rapid eye movement (REM) bout duration was more strongly associated with the amount of non-REM sleep than with intervening wakefulness. Transitions from non-REM to REM sleep were least likely when the probability of entering deep sleep peaked, in both narcolepsy type 1 and controls. Finally, transient-oscillation analysis identified a fast-sigma deficit and excess centrofrontal theta activity in narcolepsy type 1. Together, these findings connect features of sleep to disease risk, clinical classification, and the regulation of sleep itself, and demonstrate how agentic AI can support large-scale, multimodal discovery.

Introduction

Humans spend roughly one-third of their lives asleep, and disordered or insufficient sleep is consistently associated with elevated risk of cardiovascular disease, type 2 diabetes, depression, and all-cause mortality [1, 2]. Yet the biological functions of sleep remain only partially understood [3, 4, 5]. The clinical gold standard, polysomnography (PSG), captures a full night of physiology as a rich, multichannel recording [6, 7], and decades of clinical practice have accumulated vast PSG archives. Manual scoring and the high dimensionality of these recordings have nonetheless confined most studies to a handful of pre-specified hypotheses, leaving much of what these archives hold unexamined.

Sleep foundation models and related deep-learning tools have shown that a single night of PSG carries signals predictive of clinical outcomes years later, establishing that these archives are worth mining at scale [8, 9, 10, 11]. These models, though, are built to improve clinical prediction rather than to explain the physiology behind it, leaving a different class of question open. Which features within the sleep signal correlate best with which outcomes? What drives those correlations, and is it possible to move from correlation towards mechanism? Where conventional clinical categories are contested, can physiology adjudicate between them? Each of these questions requires analysing large, multimodal cohorts end to end, and it is the cost of doing so that has left the archives underused.

That cost has three sources. The first is scale: a single PSG contains hours of continuously sampled, multichannel physiology, and modern clinical archives can hold tens of thousands of such recordings [10], so a single hypothesis may require processing hundreds of thousands of hours of signal. The second is multimodal complexity: each recording integrates concurrent channels spanning the brain, autonomic, respiratory, and musculoskeletal systems, sampled at up to hundreds of hertz, each requiring distinct clinical interpretation. The third is fragmentation: recordings accumulate at different institutions under different montages and channel-naming conventions, and must be harmonized before they can be analysed together. Layered over all three is the domain expertise the data demand, spanning consensus scoring standards, sleep physiology, and the epidemiological methods needed to link one night of PSG to a diagnosis recorded years later.

Research agents built on large language models have begun to support several stages of this kind of work, including literature synthesis, hypothesis generation, code execution, data analysis, and reporting, with the shared aim of widening the range of hypotheses a scientist can examine while reducing the time spent on implementation and iterative testing [12, 13, 14, 15, 16, 17, 18, 19, 20, 21]. Some specialize in long-horizon autonomous execution from a researcher-supplied scientific objective [13, 14, 22]; others emphasize ideation and multi-agent collaboration under human direction [20, 23, 18]; others connect model reasoning to specialized scientific software and databases [12, 15, 17, 19]. These demonstrations have largely operated on literature-grounded hypothesis spaces, standard machine-learning datasets, or preprocessed tabular and matrix data rather than cohort-scale archives of raw physiological signal [13, 20, 14, 15, 22, 23].

Two existing agentic systems illustrate the gap. Biomni gives an agent access to a large library of biomedical tools, software packages, and curated databases. That design works because much of molecular biology is already standardized: single-cell analysis has agreed data formats, ready-made pipelines an agent can install and run, and reference databases that define the genes under study. Sleep physiology has not converged to the same degree. Its measurements must first be derived from raw multichannel recordings using validated detectors for sleep staging, arousals, respiratory events, and oscillatory microstructure, applied to signals whose channel names and montages differ from site to site. Deriving them is part of the scientific work rather than a step that precedes it, and Biomni offers no pipelines for overnight polysomnography, drawing on curated flat files rather than raw signal [15]. Kosmos sustains long autonomous cycles of analysis and literature search, but over supplied datasets of a few gigabytes rather than raw signal, and it fixes its research direction at the outset, leaving researchers little opportunity to intervene between cycles to steer it [14]. The archives analysed here are orders of magnitude larger, and every feature used in these case studies had to be derived from raw signal. Where scoring conventions, preprocessing, channel availability, and clinical plausibility all require expert judgement partway through an analysis, limited steerability constrains alignment between system and investigator as much as scale does. Closing this gap depends less on a more capable general-purpose agent than on the infrastructure built around one.

We therefore built AI SLEEP CO-SCIENTIST, designed around collaboration between AI agents and human sleep experts (Figure 1). A conversational Interface Layer directs three specialist agents on behalf of the expert, who uses it to formulate research questions, launch runs on the cluster, review results, and re-enter structured, file-anchored feedback. The Hypothesis Agent, seeded by a broad direction from the expert, explores data-feasible hypotheses and returns a scored portfolio to curate; the Preprocessing Agent designs

and validates feature-extraction pipelines on raw PSG before any full-cohort analysis; and the Execution Agent carries a selected hypothesis through staged statistical analysis to a report in which every reported number traces to an executable script. All components share a read-only data substrate spanning more than 50 TB of raw PSG, in which validated signal-derived features are cached and reused, so that repeated analyses of tens of thousands of recordings do not require reprocessing the underlying signals for every hypothesis.

With the support of AI SLEEP CO-SCIENTIST, we analysed roughly 124,000 PSG recordings across four cohorts through five case studies spanning three directions: how sleep physiology relates to disease risk, how it defines clinical phenotypes, and how sleep itself is organized and regulated. Two studies focused on *disease risk*: whether coupling among the brain, heart, and peripheral systems is associated with incident neurodegenerative disease, and whether multidomain sleep-age residuals reflect current disease burden, future risk, or both. One addressed *clinical phenotyping*: whether cortical arousal dynamics can reveal a clearer physiological signature of comorbid insomnia and sleep apnoea (COMISA), a contested category that conventional PSG summaries do not readily distinguish. Two examined *mechanism*: the short-term regulation of rapid eye movement (REM) sleep and its relationship to the emergence of deep sleep, and whether the transient-oscillation organization of sleep microstructure is disrupted in narcolepsy type 1 (NT1). Each study began with a broad direction proposed by a sleep researcher and was refined through exploratory rounds with AI SLEEP CO-SCIENTIST into a specific, data-feasible question before full-cohort analysis.

Across these five studies, the analyses reproduced established sleep physiology and identified patterns not previously evaluated at this scale: reduced sleep network-physiology coupling associated with incident Parkinson’s and Alzheimer’s disease diagnoses; a physiologically structured late-fusion sleep-age model whose disease associations generalized across cohorts; an obstructive sleep apnoea (OSA)-skewed intermediate COMISA phenotype characterized by prolonged post-arousal wakefulness and irregular arousal organization; a stronger association of subsequent REM-bout duration with inter-REM non-REM (NREM) duration than with intervening wake; and transient-oscillation disruptions in NT1.

AI SLEEP CO-SCIENTIST overview

AI SLEEP CO-SCIENTIST is built around two principles. First, hypothesis generation is data-driven but shaped by expert priors: the expert sets the research direction, the Hypothesis Agent explores what the available data can actually support and returns scored candidate questions, and the expert selects which of them to advance. Second, expert authority is preserved at every irreversible step, so that hypothesis selection and full-cohort commitments occur only after explicit confirmation.

A typical project is a recurring loop between the expert and the system (Figure 1a). Through a chat-based *Interface Layer* with persistent project memory, the expert brainstorms directions, launches runs, reviews reports and figures, and returns file-anchored feedback. From that interface the expert launches one of three specialist agents: the *Hypothesis Agent*, which turns a broad direction into a scored portfolio of data-feasible candidate questions; the *Preprocessing Agent*, invoked when a hypothesis requires a feature not yet extracted from raw signal; and the *Execution Agent*, which carries a chosen hypothesis through staged analysis to a code-grounded report. Each specialist coordinates a small team of sub-agents that iterate until the work passes the agent’s internal review. Expert feedback then re-enters the relevant agent as a first-class input, carrying the same authority as the original specification.

All four components operate against a shared, read-only data substrate that fixes both what data is exposed and how the agents reach it. Features are served through a schema-validated *describe*, *list*, and *load* interface rather than prompt-embedded documentation that can drift from the underlying files. A single canonical-split function returns the authoritative partition of recordings into training, validation, and test sets, so that an agent cannot invent its own split. A shared *Skills* library supplies the figure-style standards and sleep-science conventions common to every agent.

Three constraints keep the outputs grounded in the data. Candidate hypotheses must pass a data-feasibility check that rejects questions requiring variables that are neither preprocessed nor derivable from the raw signal. Feature-extraction pipelines are validated on a small sample of recordings and documented for expert review before any full-cohort extraction begins. Every reported number, figure, and table is bound to an executable script that re-derives it, with each citation resolved against its retrieved source. Across all three, an independent critic that cannot execute or alter the work it reviews audits each iteration. Supplementary Figures 7 and 8 trace two reported quantities back through the code that produced them as representative examples. The Methods section describes the Interface Layer, the three specialist agents and their sub-agents, and these safeguards in full.

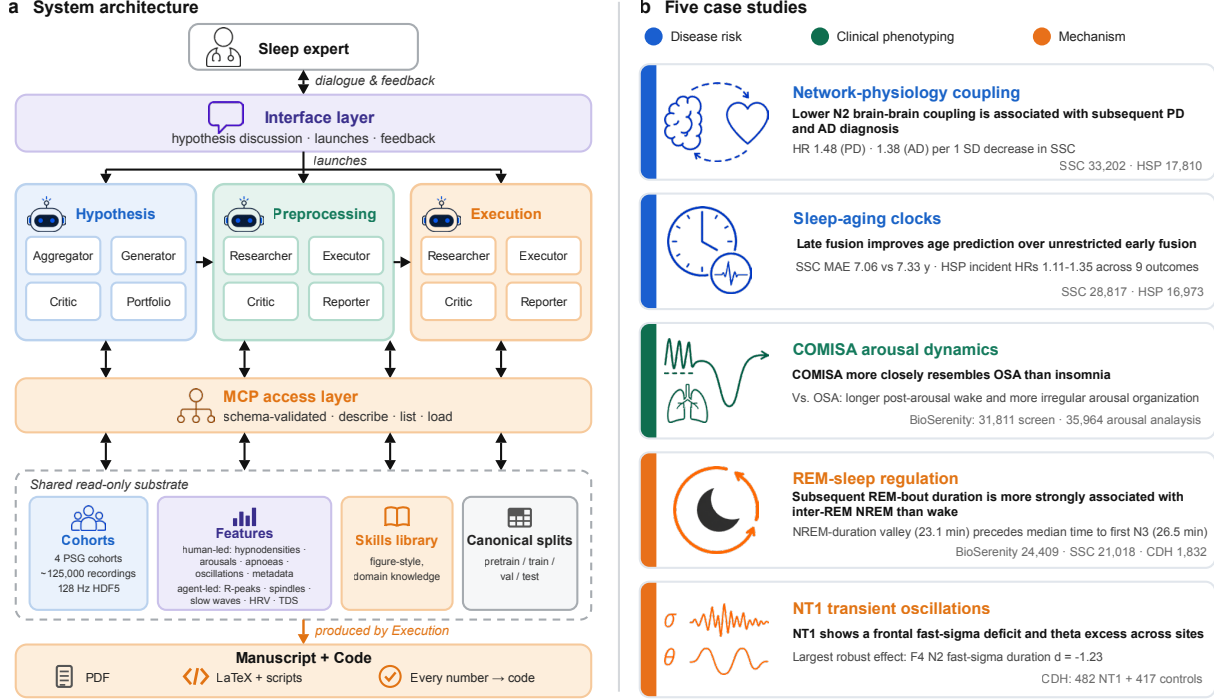


Figure 1: Overview of AI SLEEP CO-SCIENTIST. (a) System architecture. A sleep expert directs research through an *Interface Layer*, a conversational application for hypothesis discussion, agent launches, and structured feedback. The interface coordinates three specialist agents: a *Hypothesis Agent* that generates, critiques, and aggregates candidate hypotheses into a portfolio; a *Preprocessing Agent* that develops and validates physiological feature pipelines; and an *Execution Agent* that carries selected hypotheses through analysis to a code-grounded manuscript. The Hypothesis Agent comprises Aggregator, Generator, Critic, and Portfolio sub-agents, while the Preprocessing and Execution Agents each comprise Researcher, Executor, Critic, and Reporter sub-agents. The specialist agents access a shared read-only substrate through a schema-validated Model Context Protocol (MCP) layer supporting data description, listing, and loading. The substrate contains four PSG cohorts (~124,000 recordings; > 50 TB), curated physiological features, a shared skills library, and canonical train/validation/test splits. The Execution Agent returns a manuscript and the associated $\text{\LaTeX}$ files, analysis scripts, and executable provenance. (b) Five case studies spanning disease risk, clinical phenotyping, and mechanism, each summarized by its principal finding, supporting result, and analysis cohorts.

Data and research questions

AI SLEEP CO-SCIENTIST drew on four clinical and epidemiological PSG cohorts: the Stanford Sleep Clinic (SSC) [8, 10], the Human Sleep Project (HSP) [24], the BioSerenity clinical network [25, 26, 27], and the Central Disorders of Hypersomnolence (CDH) registry [28, 29], together comprising roughly 124,000 PSG recordings (Supplementary Table 2). They differ in design, and these differences determine which questions each can address. SSC and HSP are large clinical PSG archives with linked electronic health records providing longitudinal follow-up and diagnoses coded with the International Classification of Diseases (ICD). They are the only cohorts in the substrate that support prospective analyses linking a single night of sleep to a diagnosis recorded years later, and because they were assembled independently at different institutions, SSC can serve as a discovery cohort and HSP as external validation. BioSerenity is a multi-site clinical network of comparable scale that lacks linked diagnostic codes but contains sleep questionnaires, respiratory measures, and symptom histories, making it suited to phenotyping questions whose group definitions rest on symptoms combined with physiology rather than on coded diagnoses. CDH is a smaller multi-site registry with harmonized diagnoses of the central disorders of hypersomnolence, including narcolepsy type 1, and therefore supports mechanistic questions in a population with a specific and well-characterized neurochemical lesion.

To support rapid hypothesis execution, we precomputed a curated feature library across the four cohorts, spanning sleep macrostructure and microstructure together with autonomic, respiratory, and brain-physiological coupling (Supplementary Table 1). The human-led features derive from validated community tools, including sleep staging from U-Sleep [30], cortical arousals from the Multimodal Arousal Detector [31], apnoea events from the Apnoea-Based Event Detector [32], transient oscillations from DynamO [33], and an established formulation of brain–heart coupling [34], alongside each cohort’s curated clinical metadata. The agent-led features were curated, adapted, or newly implemented by the Preprocessing Agent under expert review. Some are adaptations of existing community tools that the agent identified from the literature, pulled from source repositories, adapted for compatibility across cohorts, and exposed through the MCP interface. These include event-level microstructure such as spindles, slow waves, eye movements, and EEG artefact masks detected with Yet Another Spindle Algorithm (YASA) [35], and stage-stratified physiological coupling networks quantified by time-delay stability (TDS) [36]. Others were designed by the agent directly from published methods, including R-peak detection, continuous oximetry, desaturation events, and heart-rate-variability (HRV) summaries [37]. Building on the original brain–heart coupling formulation, the agent further constructed a brain–EMG coupling measure on the same principle [34]. All features were harmonized and post-hoc verified across cohorts.

The five areas of investigation were not fixed in advance of the data. Each began as a broad direction proposed by a sleep researcher from an open question in the field, was checked against what the substrate could actually support, and was then sharpened into a specific question through exploratory rounds between the investigator and AI SLEEP CO-SCIENTIST before any full-cohort analysis was committed. The Hypothesis Agent triaged candidate questions for scientific merit and data feasibility, and the investigator used the Interface Layer to select which of them advanced (Figure 1b). The five case studies therefore reflect the intersection of expert research priorities with the particular strengths of the available cohorts.

Although the five studies address different disorders and timescales, they share a common setting: each asks what a single night of polysomnography reveals at a particular level of physiological organization. Those levels differ across the studies, spanning coupling between physiological systems, the domain composition of physiological ageing, the temporal organization of cortical arousals, the duration structure of cycling between NREM stages 1–3 (N1–N3) and REM sleep, and oscillatory events shorter than a scoring epoch. These questions group into three directions. Two case studies concern *disease risk*, motivated by the observation that a single night of sleep carries signal predictive of disease years later without a shared framework for determining which features carry which risks. They ask whether coupling among brain, heart, and peripheral systems is associated with the risk of incident neurodegenerative disease, and whether multidomain sleep-age residuals are associated separately with prevalent disease burden and future disease risk. One concerns *clinical phenotyping*, motivated by contested clinical classifications that conventional PSG summaries cannot resolve: whether the dynamics of cortical arousal can decompose COMISA into a clearer physiological signature. Two concern *mechanism*, motivated by unresolved questions about how sleep is biologically organized: whether short-term REM-sleep regulation tracks preceding NREM sleep, how inter-REM NREM duration is distributed in relation to N3 emergence across the night, and whether the transient-oscillation organization of sleep microstructure is disrupted in NT1. The remainder of this paper presents the findings from each case study in turn.

Disorder-specific network-physiology coupling profiles during sleep

Main finding

Neurological and neuropsychiatric disorders, particularly Parkinson’s disease and Alzheimer’s disease, showed distinct sleep-stage- and subnetwork-specific alterations in physiological coupling with moderate concordance across SSC and HSP. In SSC, each 1 SD decrease in N2 brain–brain coupling was associated with a higher risk of subsequent Parkinson’s disease (HR 1.48, 95% CI 1.31–1.67) and Alzheimer’s disease (HR 1.38, 95% CI 1.25–1.53) diagnoses.

Sleep is a coordinated physiological state in which interactions between the brain and body continually reorganize across sleep stages. Neurological and neuropsychiatric disorders may therefore alter not only individual PSG signals, but also how different physiological systems interact over time. Consistent with this possibility, a recent foundation-model study of polysomnography found that cross-modal representations learned across electroencephalographic, respiratory, and other physiological signals outperformed single-modality representations on health-prediction tasks [10]. However, it remains unclear whether different




Figure 2: Network-physiology coupling during sleep across neurological and neuropsychiatric disorders. (a) Study design. EEG and ECG-derived heart-rate signals from overnight polysomnography were analysed within Wake, N1, N2, N3, and REM sleep using time-delay stability (TDS). TDS applies sliding-window cross-correlation to pairs of physiological signals and retains links for which the estimated time delay remains stable in more than 7% of windows, producing subject-level brain-brain and brain-heart coupling networks for each sleep stage. Cross-sectional effect sizes were estimated in SSC (discovery; controls, $n = 11,240$; eight disorder groups) and replicated in HSP (controls, $n = 3,604$; the corresponding eight disorder groups). (b) Brain-brain and brain-heart TDS link strength across sleep stages in SSC controls and individuals with Parkinson's disease or Alzheimer's disease. Points and error bars show means and 95% confidence intervals. (c) Brain-heart TDS coupling in SSC controls, stratified by sleep stage and EEG frequency band (δ , θ , α , σ , and β). (d) Cross-site concordance of covariate-adjusted TDS effect sizes, shown as SSC versus HSP Hedges' g across 2,496 disorder-feature pairs (Pearson $r = 0.631$). Positive values indicate lower coupling in the disorder group than in controls, whereas negative values indicate higher coupling. Points in the upper-right and lower-left quadrants therefore show concordant effect directions across cohorts. (e) Prospective analysis in SSC. Points and error bars show hazard ratios and 95% confidence intervals from Cox proportional-hazards models per 1 SD decrease in brain-brain or brain-heart coupling, stratified by incident outcome and whole-night or sleep-stage-specific coupling summary. Prevalent cases were excluded separately for each outcome. Open markers denote associations with $q > 0.05$ after false-discovery-rate (FDR) adjustment.

disorders produce distinct alterations in physiological coupling during sleep, whether these patterns generalize across cohorts, and whether disrupted coupling precedes subsequent disease diagnosis.

To characterize these interactions, we used time-delay stability, a framework for identifying stable coupling between physiological signals [36]. TDS calculates the cross-correlation between each pair of signals across successive sliding windows and identifies the time lag at which the correlation is maximal. Link strength was defined as the proportion of windows in which this lag remained stable, with a signal pair classified as significantly coupled when this proportion exceeded 7%. These pairwise measurements were used to construct sleep-stage-specific physiological networks.

The analysis used SSC as the discovery cohort ($n = 33,202$) and HSP as an independent replication cohort ($n = 17,810$) [24]; cohort selection and group counts are reported in Supplementary Table 3. In each cohort, the Control group was identified using PheCode-based criteria and was required to have no recorded history of neurological disorders (epilepsy, Parkinson’s disease, dementia, Alzheimer’s disease, or mild cognitive impairment [MCI]); neuropsychiatric disorders (major depressive disorder [MDD] or post-traumatic stress disorder [PTSD]); cardiovascular diseases (hypertension, atrial fibrillation, heart failure, acute myocardial infarction, or ischaemic heart disease); sleep disorders (insomnia or REM sleep behaviour disorder); or cerebrovascular events, including stroke or transient ischaemic attack. The Control groups comprised 11,240 SSC and 3,604 HSP participants. These participants were compared with eight neurological and neuropsychiatric disorder groups to determine whether different disorders exhibited shared or distinct coupling profiles and whether reduced coupling was associated with subsequent disease diagnosis.

The Control group provided a baseline for characterizing network physiology within each sleep stage. Within SSC, the brain–brain and brain–heart subnetworks showed distinct sleep-stage-dependent profiles. Brain–brain coupling increased from Wake to N1 (link strength: 24.1% and 31.9%, respectively), peaked during N2, declined during N3, and reached its lowest level during REM sleep. Brain–heart coupling followed a different profile, peaking during N1 before decreasing through N2 and reaching its minimum during N3 (link strength: 3.5%) (Figure 2b,c). Global brain–body coupling was also strongest during N2 and lowest during REM sleep. The attenuation of coupling from lighter NREM sleep to N3 was consistent with previous TDS findings [36]. The feature landscapes of the SSC and HSP Control groups were highly correlated across 70 features ($r = 0.979$, $p < 10^{-48}$). Nevertheless, differences in absolute coupling remained between cohorts: brain–heart features showed small standardized differences, whereas brain–brain coupling was systematically higher in SSC than in HSP, particularly during N2 and N3.

To characterize disorder-specific network-physiology profiles, the Control group was compared separately with PheCode-defined groups for epilepsy, partial epilepsy, Parkinson’s disease, Alzheimer’s disease, dementia, MCI, PTSD, and MDD (Supplementary Figures 9 and 10). Analyses were conducted separately in SSC and HSP, with SSC serving as the discovery cohort and HSP as the replication cohort. Cross-sectional differences were quantified using Hedges’ g after adjustment for age, sex, and body mass index (BMI) in SSC and for age and sex in HSP, where positive values of g indicated lower coupling in the disorder group than in the Control group.

In SSC, Parkinson’s disease showed the broadest reduction in network coupling, with the largest effect observed for the percentage of significant frontal brain–brain links during N1 (Hedges’ $g = 0.86$). Alzheimer’s disease, dementia, and MCI showed similar reductions in frontal–central brain–brain link strength during N1 (Hedges’ $g = 0.68$, 0.58 , and 0.51 , respectively). Epilepsy was associated with modestly lower frontal brain–brain coupling during N1 and N2, with the larger effect occurring during N1. Partial epilepsy showed its largest effect in the percentage of significant frontal brain–brain links during N2 (Hedges’ $g = 0.51$).

The neuropsychiatric groups showed more selective profiles. PTSD had the lowest proportion of FDR-significant features (23.4%) and modest effect sizes (Hedges’ $|g| = 0.30$), with its largest effect observed for a reduction in the percentage of significant brain–EMG links during N2. In MDD, the largest overall effect was a reduction in brain–EOG coupling during N2. However, MDD differed from the other groups in showing predominantly higher brain–heart coupling relative to controls.

Cross-site replication was assessed by repeating the disorder comparisons in HSP and comparing the resulting effect-size patterns with those observed in SSC. Across 2,496 disorder–feature pairs, Hedges’ g values showed moderate concordance between cohorts (Pearson $r = 0.631$; Figure 2d). Concordance was strongest for Parkinson’s disease, dementia, and MDD, and weakest for PTSD. Because BMI was unavailable in HSP, the SSC analysis was repeated with and without BMI adjustment for five representative features across all eight disorder groups. The largest observed difference was for the percentage of significant frontal brain–brain links during N1 in Parkinson’s disease (Hedges’ $g_{\text{with BMI}} = 0.877$; Hedges’ $g_{\text{without BMI}} = 0.908$). Adjustment for

BMI did not change the disorder rankings or which comparisons were statistically significant, indicating that BMI adjustment was unlikely to account for the cross-site agreement observed for the tested features.

To determine whether reduced coupling was associated with subsequent diagnoses, we fitted Cox proportional-hazards models to time-to-event data in SSC and HSP. For each outcome, participants with a diagnosis recorded before the PSG were excluded, and follow-up extended from the PSG date to the first recorded diagnosis or the last documented clinical contact. Models were adjusted for age, sex, and BMI in SSC and for age and sex in HSP. Hazard ratios estimated the change in risk associated with each 1 SD decrease in coupling. In SSC, lower brain–brain coupling was associated with an increased risk of all seven incident outcomes examined, with the strongest associations generally observed during N2 sleep (Figure 2e).

Lower brain–heart coupling was also associated with a higher risk of subsequent Parkinson’s disease, Alzheimer’s disease, and dementia diagnoses, primarily during N2, N3, and REM sleep. The association with Parkinson’s disease was particularly pronounced (Figure 2b,e). This finding is biologically plausible because progressive loss of cardiac noradrenergic sympathetic innervation is a characteristic feature of Parkinson’s disease and can occur independently of nigrostriatal dopaminergic degeneration [38].

For N2 brain–brain coupling, each 1 SD decrease was associated with a higher risk of subsequent Parkinson’s disease (HR 1.48, 95% CI 1.31–1.67) and Alzheimer’s disease (HR 1.38, 95% CI 1.25–1.53) diagnoses (Figure 2e). The observed reduction in brain–brain coupling in these two neurodegenerative disorders is biologically plausible. One possible explanation is that neurotransmitter-specific deficits in dopamine in Parkinson’s disease and acetylcholine in Alzheimer’s disease, together with proteinopathy-related synaptic and interneuron impairment, disrupt the excitatory–inhibitory balance and structural connectivity required for coordinated oscillatory coupling across brain networks.

Expert note on collaboration

The investigator framed the initial question based on previous work on network-physiology coupling during sleep, and exploratory hypotheses were subsequently refined with AI SLEEP CO-SCIENTIST in response to the primary findings. For each analytical step, the investigator specified the objective, required outputs, and scope constraints. Within these constraints, AI SLEEP CO-SCIENTIST autonomously implemented the PheCode-based group definitions, extracted and summarized features from the precomputed TDSpy outputs, performed the effect-size analyses, assessed cross-site replication, fitted the Cox survival models, and generated the figures. Across four iterative rounds, feedback from the investigator and domain expert focused on clarifying the methodology, narrowing the analytical scope, and decomposing complex requests into sequential, single-objective tasks. AI SLEEP CO-SCIENTIST incorporated this feedback into subsequent analyses. This process also revealed that broad, multi-part prompts could lead AI SLEEP CO-SCIENTIST to prioritize the most prominent findings while overlooking other elements of the requested scope, motivating the use of narrower prompts with explicitly defined outputs.

Multidomain polysomnographic sleep ageing reflects current disease burden and future disease risk

Main finding

Physiologically structured late fusion outperformed unrestricted early fusion for PSG-based age prediction in the Stanford Sleep Clinic cohort and generalized more accurately to the Human Sleep Project cohort. Its bias-corrected age residual was associated with current disease burden and a reproducible pattern of incident cardiometabolic, renal, and respiratory disease. Disease-specific reweighting of the seven domain residuals did not improve held-out discrimination of the broad incident-disease composite.

Sleep-derived brain-age models can predict chronological age from PSG, and differences between predicted and chronological age have been associated with mortality, dementia, and other clinical outcomes [39, 40, 41]. One of the most widely used measures, the Brain Age Index, showed a mean absolute deviation of 7.6 years from chronological age among healthy participants in the Massachusetts General Hospital dataset and has also been associated with crystallized cognition [42]. Yet these associations do not establish what an elevated age residual represents: a general process of accelerated biological ageing, vulnerability within a particular physiological system, or changes resulting from established illness. Although shared mechanisms of ageing may increase susceptibility to disease, individual disorders also arise from organ-specific pathology, genetic predisposition, environmental exposures, and the consequences of existing illness [43]. We therefore asked

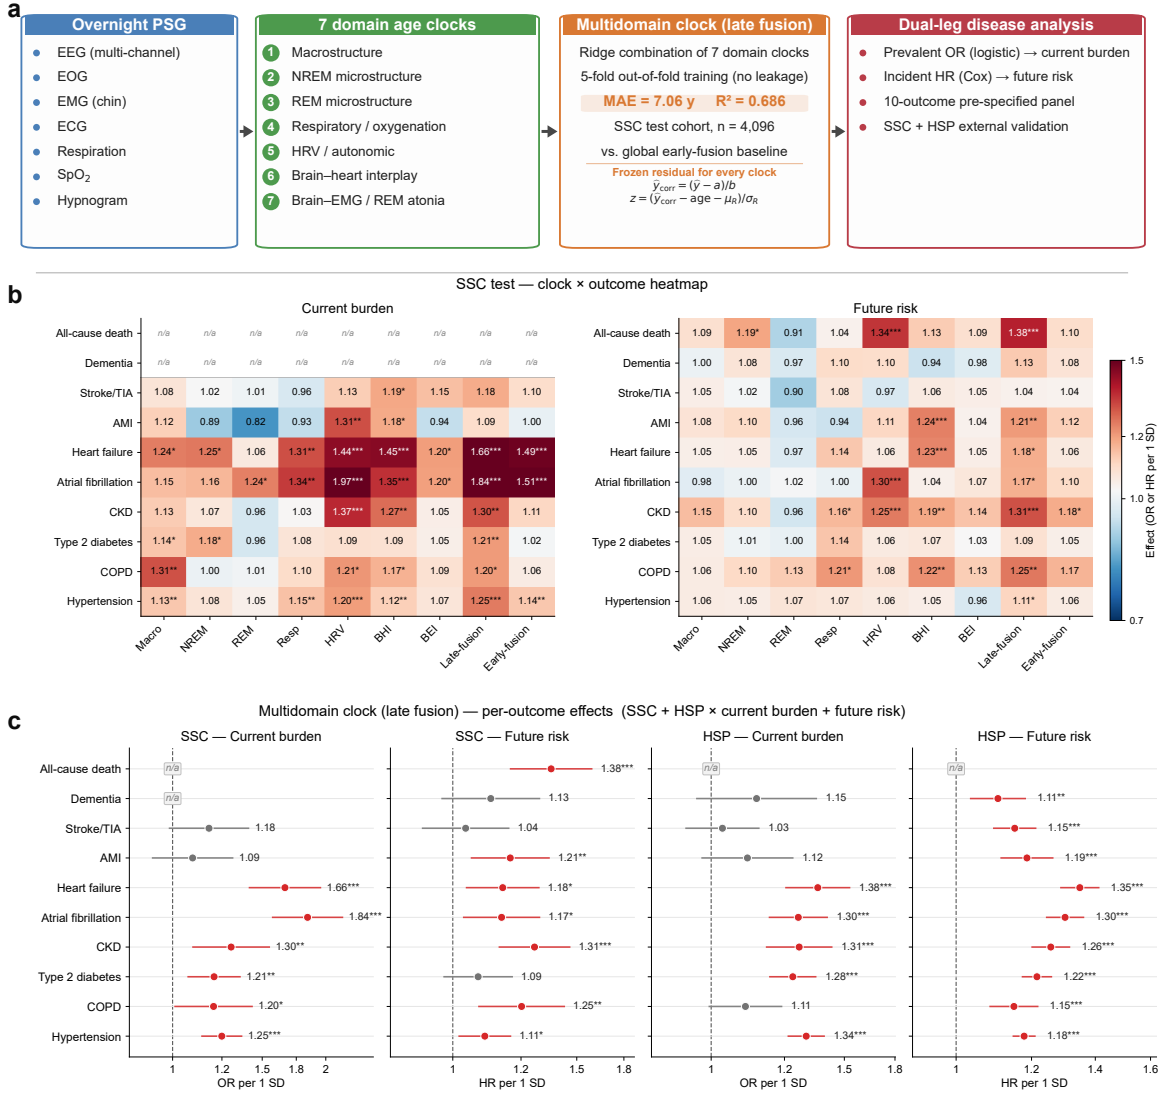


Figure 3: Multidomain polysomnographic sleep-ageing residuals are associated with prevalent disease burden and incident disease risk. (a) Analysis pipeline. PSG signals were represented by seven physiologically interpretable domain age clocks: macrostructure, NREM microstructure, REM microstructure, respiratory/oxygenation physiology, HRV/autonomic physiology, brain-heart interplay, and brain-EMG/REM-atonia. Predictions from the seven clocks were combined using a ridge late-fusion model trained on five-fold out-of-fold predictions. The resulting model achieved a mean absolute error (MAE) of 7.06 years and $R^2 = 0.686$ in the SSC test set ($n = 4,096$ of the 4,120 test-split recordings, the remainder lacking required features; Supplementary Table 4). Predicted ages were linearly bias-corrected, and residuals were standardized using parameters estimated from SSC training and validation recordings with an apnoea-hypopnoea index (AHI) ≤ 50 . These parameters were then fixed and applied to the SSC test set and HSP. (b) Clock-by-outcome heatmaps for the prespecified 10-outcome panel in SSC, shown for the seven domain clocks, the multidomain late-fusion model, and the global early-fusion baseline. Prevalent associations were estimated using logistic regression and represent disease burden at the time of PSG. Incident associations were estimated using Cox models after excluding prevalent cases for each outcome and represent future disease risk. Effects are reported per 1 SD increase in the corresponding age residual. (c) Associations of the multidomain late-fusion residual with individual outcomes in SSC and HSP. Points and lines show effect estimates and 95% confidence intervals; red denotes nominal significance and grey denotes $p \geq 0.05$. Asterisks denote $*p < 0.05$, $**p < 0.01$, and $***p < 0.001$. n/a indicates that prevalent mortality was not applicable, mortality follow-up was unavailable in HSP, or the prevalent-dementia model in SSC did not converge because only 27 cases were available. Incident associations were broadly consistent across outcomes available in both cohorts.

whether PSG-derived age residuals reflect a single global ageing axis or heterogeneous ageing trajectories across physiological systems, and whether their associations with disease capture existing disease burden, predict future disease risk, or both.

The analysis was developed in the SSC cohort and externally validated in a Harvard subsample of the Human Sleep Project cohort [24]. After harmonized quality control, the analysis included 28,817 SSC recordings and 16,973 HSP recordings (Supplementary Table 4). Seven domain-specific clocks captured sleep macrostructure, NREM and REM microstructure, respiratory and oxygenation physiology, HRV and autonomic physiology, brain–heart interplay, and brain–EMG and REM-atonía physiology. A ridge late-fusion model combined five-fold out-of-fold predictions from these clocks and achieved the best age-prediction performance on the held-out SSC test set, with an MAE of 7.06 years and $R^2 = 0.686$. By comparison, the strongest single-domain clock achieved an MAE of 8.05 years, and the unrestricted global early-fusion model achieved an MAE of 7.33 years (Figure 3a; Supplementary Figure 11a). Late fusion reduced MAE relative to early fusion by 0.27 years in paired-bootstrap analysis (95% CI, 0.13–0.42 years), supporting the value of combining predictions across physiologically defined domains rather than directly aggregating all features. For each clock, predicted ages were corrected for age-related bias, and the resulting age residuals were standardized using parameters estimated from SSC training and validation recordings with $AHI \leq 50$. All model and correction parameters were then frozen before evaluation in the SSC test set and external validation in HSP.

The principal disease analysis evaluated a prespecified panel of 10 outcomes in parallel prevalent and incident analyses, where applicable. The panel included all-cause mortality, dementia, cardiovascular diseases, chronic kidney disease (CKD), type 2 diabetes, and chronic obstructive pulmonary disease (COPD). Logistic regression was used to test whether the age residual was associated with disease already present at the time of PSG. For each incident analysis, participants with the corresponding disease at baseline were excluded, and Cox regression was used to assess subsequent disease risk. This parallel design allowed each outcome to be interpreted as reflecting current disease burden, future disease risk, both, or neither.

The associations did not support a uniform, global signal in which sleep-age residuals predicted all outcomes. Instead, they followed a coherent pattern across cardiometabolic, renal, and respiratory disease (Figure 3b). The late-fusion and HRV residuals showed the broadest associations, whereas the brain–heart interplay residual was particularly associated with cardiovascular outcomes. In incident analyses, the late-fusion residual was associated with all-cause mortality, CKD, COPD, acute myocardial infarction, heart failure, atrial fibrillation, and hypertension; associations with dementia and stroke or transient ischaemic attack were weaker in SSC. Prevalent associations were strongest for heart failure, atrial fibrillation, CKD, type 2 diabetes, COPD, and hypertension. Together, these patterns suggest that the residual reflects a combination of shared physiological vulnerability, disease-specific dysfunction, and changes associated with established disease, rather than a single causal pathway of ageing.

External validation provided further support for the physiologically structured late-fusion approach. When applied to HSP without retraining, age-prediction performance declined, potentially reflecting cohort shift; nevertheless, late fusion generalized better (MAE 8.83 years, $R^2 = 0.494$) than the global early-fusion baseline (MAE 9.52 years, $R^2 = 0.315$). Associations between the late-fusion residual and individual outcomes were also broadly consistent across SSC and HSP (Figure 3c). Incident associations were particularly consistent, whereas prevalent associations varied more between cohorts, possibly owing to differences in referral patterns, baseline disease burden, clinical coding, and covariate availability. Disease-tuned models were evaluated as secondary risk scores rather than as alternative ageing clocks. For the broad disease composite, disease-specific reweighting of the domain residuals provided no improvement in discrimination over the age-derived late-fusion residual in either SSC ($\Delta C = -0.005$ vs. $+0.002$) or HSP ($\Delta C = +0.004$ vs. $+0.005$), relative to a demographic baseline model.

A phenome-wide association study (PheWAS) extended these findings beyond the prespecified 10-outcome panel (Supplementary Figure 11b). In SSC, the late-fusion residual was associated at FDR significance with 39 incident and 80 prevalent PheCodes. Associations were concentrated in disease categories related to the primary panel, including circulatory, endocrine/metabolic, renal, genitourinary, and respiratory disorders. Among the SSC-significant PheCodes that could be evaluated in HSP, effect directions were concordant for 100% of incident and 94% of prevalent associations. This breadth suggests that the multidomain PSG age residual captures clinically meaningful physiological dysfunction, whereas its concentration within particular organ systems argues against interpreting it as a universal mediator of disease.

To assess whether the findings depended on a particular agent trajectory, the complete study was independently repeated three times from the same high-level prompt (Supplementary Table 5). Although the agent developed different feature spaces across runs, leading to variation in age-prediction accuracy, the main findings

were consistently reproduced: associations spanning cardiometabolic, renal, respiratory, and autonomic disease; generalization of incident PheWAS associations; and limited incremental benefit from disease-specific reweighting of the age residuals.

The value of this case study lay not only in the model development but in how it refined the scientific question. Rather than treating a PSG-derived age gap as a universal measure of biological ageing [39, 40, 41, 42], the analysis decomposed sleep ageing into physiologically defined domains and showed that late fusion yielded a more interpretable representation that generalized better across cohorts. The resulting residual does not have a single biological meaning: depending on the outcome, it may reflect shared age-related vulnerability, disease-specific pathophysiology, consequences of established disease, or a combination of these processes [43]. Relative to previous sleep-age studies, the principal advance lies in combining physiologically structured multimodal age modelling with parallel evaluation of prevalent disease burden and incident disease risk, followed by external validation. Together, these analyses revealed a reproducible pattern of associations centred on cardiometabolic, renal, respiratory, and autonomic dysfunction, rather than a universal signal across all diseases.

Expert note on collaboration

The investigator established the initial direction based on prior literature, and the study was subsequently expanded to compare multiple global modelling strategies as the agent’s exploratory runs surfaced additional possibilities. Before scaling the analysis, the investigator specified the seven physiologically interpretable feature domains, the parallel prevalent and incident disease analyses, the prespecified 10-outcome clinical panel, and the SSC-to-HSP external-validation design. Within this framework, AI SLEEP CO-SCIENTIST engineered the features for each domain; trained the seven domain clocks, the multidomain late-fusion clock, and the global early-fusion baseline; conducted both disease analyses; and generated all figures. Methodological contributions flowed in both directions: AI SLEEP CO-SCIENTIST proposed the age bias-correction procedure that was subsequently adopted into the pipeline and independently initiated sensitivity analyses designed to challenge the main findings. Across three rounds of structured feedback through the Interface Layer, the investigator requested additional features, clearer reporting of data exclusions, and figure-level revisions. The agent incorporated each request and, where the results warranted further investigation, added further sensitivity analyses.

Arousal dynamics reveal an obstructive sleep apnoea-skewed intermediate phenotype in COMISA

Main finding

Within the sleep-microstructure feature space examined here, comorbid insomnia and sleep apnoea (COMISA) more closely resembled obstructive sleep apnoea (OSA) than insomnia. Relative to OSA, COMISA was characterized by longer post-arousal wakefulness and greater within-night irregularity in arousal organization rather than greater respiratory disturbance. This suggests that, although respiratory events may trigger cortical arousals similarly in OSA and COMISA, individuals with COMISA have greater difficulty re-establishing stable sleep following arousal. Insomnia showed minimal separation from the Control group. Together, these findings identify cortical arousal dynamics as a complementary dimension of COMISA physiology beyond conventional respiratory measures.

Obstructive sleep apnoea and insomnia are among the most common sleep disorders and frequently co-occur [44]. This combined phenotype, termed comorbid insomnia and sleep apnoea, affects an estimated 30–50% of individuals with either disorder and is associated with greater daytime impairment, elevated cardiometabolic risk, and poorer treatment outcomes than either condition alone [45, 46, 47]. Despite its clinical importance, the nocturnal physiology underlying COMISA remains poorly understood, with prior studies limited to small samples [48, 49]. A dominant model conceptualizes COMISA as a hybrid disorder arising from the interaction between insomnia-related hyperarousal and OSA-driven respiratory instability [48]. However, it remains unclear whether COMISA represents a balanced integration of these mechanisms or is physiologically dominated by one constituent process.

There remains debate over whether COMISA is characterized by a reduced arousal threshold, meaning a greater propensity to awaken in response to an apnoeic event [49, 50]. Conventional PSG metrics, including sleep efficiency, sleep-stage composition, and the AHI, provide only coarse summaries of sleep and respiration




Figure 4: Arousal dynamics characterize COMISA as an OSA-skewed intermediate phenotype. (a) Case-study schematic. Clinical polysomnography from the BioSerenity cohort was analysed across four diagnostic groups (Control, insomnia, OSA, COMISA; inset definitions). Four sleep-microstructure domains were first screened on $n_1 = 31,811$ recordings, followed by arousal dynamics with 46 features across four physiological families on $n_2 = 35,964$ recordings. Pairwise logistic classification revealed differences in diagnostic separability, with some comparisons classified with high accuracy (area under the receiver operating characteristic curve [AUC] ≈ 0.95) and others proving substantially more challenging (AUC ≈ 0.65). Three of four principal claims replicated across the SSC and HSP cohorts. (b) Discrimination summary. *Left*: macro-averaged one-vs-rest AUC per single sleep-microstructure domain (hypnodensity, arousals, spindles, slow waves); cortical arousals are the most informative single domain (0.765 versus 0.744–0.754 for the others). *Right*: pairwise arousal-dynamics AUC with 95% CIs: OSA versus Control (0.957 [0.951, 0.962]) and COMISA versus insomnia (0.950 [0.945, 0.956]) are near-ceiling, whereas COMISA versus OSA (0.651 [0.641, 0.666]) and insomnia versus Control (0.634 [0.614, 0.652]) are considerably harder. (c) Covariate-adjusted standardized coefficients distinguishing COMISA from OSA, coloured by arousal-feature domain; positive values denote features elevated in COMISA. Longer post-arousal wakefulness and more irregular arousal organization (local arousal-density coefficient of variation) are the strongest COMISA-leaning signals. (d) Distribution of a hybrid score (0 = insomnia-like, 1 = OSA-like) for each group; the COMISA distribution is shifted towards the OSA end (mean 0.537; $p < 10^{-300}$ versus 0.5), comprising 26.9% OSA-like, 63.4% intermediate, and 9.7% insomnia-like recordings.

and are ill suited to addressing this question [27]. A recent meta-analysis of 16 studies found that patients with COMISA had a higher microarousal index, higher AHI, lower minimum SpO₂, longer wake after sleep onset, and lower sleep efficiency than patients with OSA alone [51]. Patients with COMISA also showed longer awakenings (≥ 5 minutes), with significantly more wake time attributable to these extended awakenings than patients with OSA alone (43.3 versus 25.5 minutes, $p < 0.001$) [51]. Multi-night home monitoring further showed that patients with COMISA had greater night-to-night variability in sleep-onset latency, sleep efficiency, and awakening duration, a pattern that cannot be captured by single-night PSG [51]. Together, these limitations make conventional PSG metrics poorly suited to capturing the temporal dynamics of cortical arousals, which mediate interactions among respiratory events, autonomic activation, and sleep fragmentation [52].

To investigate this question, two analytic samples were derived from the BioSerenity PSG cohort (Supplementary Table 6). We first screened four sleep-microstructure domains, namely cortical arousals, sleep spindles, slow waves, and hypnodelta-derived sleep-stage uncertainty, using 31,811 PSG recordings from the Control, insomnia, OSA, and COMISA groups (Figure 4a). Two findings emerged. First, insomnia showed limited objective separation from the Control group, consistent with previous reports of limited reproducible PSG abnormalities in chronic insomnia [53, 54]. Second, cortical arousals were the most informative feature domain, with measures of arousal frequency, timing, and organization outperforming spindle-, slow-wave-, and staging-derived metrics in distinguishing among the diagnostic groups. Cortical arousals achieved the highest single-domain discrimination, with a macro-averaged AUC of 0.765, compared with 0.744–0.754 for the other domains (Figure 4b; Supplementary Figure 12a,c).

We then analysed a larger set of 35,964 PSG recordings for which reliable arousal detection was available, extracting 46 features describing arousal burden, respiratory coupling, temporal clustering, and post-arousal sleep–wake trajectories. Respiratory-linked arousal dynamics strongly distinguished groups that differed in sleep-disordered breathing. Pairwise logistic-regression models distinguished OSA from the Control group with an AUC of 0.957 (95% CI 0.951–0.962) and COMISA from insomnia with an AUC of 0.950 (95% CI 0.945–0.956) (Figure 4b). In contrast, COMISA was substantially more difficult to distinguish from OSA (AUC 0.651, 95% CI 0.641–0.666), as was insomnia from the Control group (AUC 0.634, 95% CI 0.614–0.652). Nevertheless, the differences between COMISA and OSA were concentrated in arousal features beyond conventional measures of respiratory-event burden. The strongest discriminative features captured post-arousal wakefulness and the temporal organization of arousal events across the night (Figure 4c). Relative to OSA alone, COMISA was characterized by longer wakefulness after arousal and greater irregularity in arousal timing. These findings are consistent with a recent study of 72 patients with COMISA and 72 patients with OSA, which found greater autonomic responses and longer arousals in COMISA [55].

Arousal dynamics also provided information beyond conventional sleep and respiratory measures: adding these features to standard PSG indices improved diagnostic discrimination across all comparisons (Δ AUC range: +0.042–0.123). To characterize the position of COMISA along the insomnia–OSA arousal spectrum, we computed a Mahalanobis-based hybrid score using the 15 features with the largest group effects. Features were z-scored, and the insomnia and OSA group centroids were estimated using a regularized pooled covariance matrix. Each recording’s score was defined by its relative Mahalanobis distance to the two centroids, ranging from 0 (insomnia-like) to 1 (OSA-like). COMISA recordings formed a largely continuous distribution rather than separating into discrete insomnia- and OSA-dominant clusters (Figure 4d; Supplementary Figure 12b). The distribution was shifted towards the OSA end of the spectrum (mean hybrid score: 0.537; $p < 0.001$ relative to the midpoint of 0.5), with 26.9% of recordings classified as OSA-like, 63.4% as intermediate, and 9.7% as insomnia-like. These findings do not contradict previous reports that insomnia and OSA can interact synergistically to affect conventional respiratory measures and clinical outcomes [56]. Rather, they indicate that, within the sleep-microstructure feature space examined here, COMISA was more closely aligned with OSA than with insomnia, while its differences from OSA were concentrated in post-arousal wakefulness and the temporal organization of arousals.

To assess generalizability, the domain-screening and arousal-focused analyses were repeated in the independent SSC and HSP cohorts using the same diagnostic definitions and feature-extraction pipelines (Figure 12). Three findings were consistent across cohorts: insomnia showed minimal separation from the Control group, cortical arousals remained the most informative microstructure domain, and the residual distinction between COMISA and OSA was concentrated in post-arousal wakefulness rather than respiratory burden. This pattern is also consistent with a smaller study reporting increased cardiorespiratory responses to apnoeas during “subcortical arousals” in COMISA compared with OSA [55]. The precise positioning of COMISA along the insomnia–OSA spectrum was more cohort-dependent. BioSerenity placed COMISA at an OSA-skewed intermediate position, SSC found it largely indistinguishable from OSA, and HSP showed greater phenotypic polarization and

evidence of sleep disruption exceeding that observed in either constituent disorder. The HSP finding is consistent with a previous large-scale study showing that insomnia and OSA can interact synergistically to amplify clinical symptoms and selected conventional PSG abnormalities, including respiratory-disturbance and arousal-related measures [56]. Respiratory-axis classification was also weaker in the replication cohorts than in BioSerenity. Despite these quantitative differences, all three cohorts supported the same broad interpretation: within the sleep-microstructure feature space examined here, COMISA was generally more closely aligned with OSA than with insomnia, although synergistic interactions between the two disorders may emerge in other physiological domains.

Together, these findings show that conventional PSG metrics do not fully capture the physiological heterogeneity of COMISA. Measures derived from cortical arousals provide complementary information and may improve physiological stratification beyond standard respiratory indices.

Expert note on collaboration

The investigator framed the initial direction based on prior work on the insomnia–apnoea overlap, while AI SLEEP Co-SCIENTIST surveyed the available physiological features and cohorts and proposed candidate domains through which the question could be addressed. The investigator then fixed the scientific question, the four diagnostic labels (Control, insomnia, OSA, and COMISA), and the training, validation, and test partitions. Within these constraints, AI SLEEP Co-SCIENTIST designed and executed the computational analyses, selected the modelling approaches, evaluated the features, performed cross-cohort validation, and generated all figures. Across four iterative rounds, investigator feedback refined the sub-hypotheses, clarified the analytical scope, and focused subsequent analyses on arousal dynamics as the most informative domain. This process also showed that shorter prompts with a single objective consistently produced more focused and reproducible outputs than compound, multi-part prompts.

REM-bout duration is more strongly associated with NREM than wake, and a valley in inter-REM NREM duration precedes N3 emergence

Main finding

Within inter-REM intervals, REM-bout duration was more strongly associated with NREM duration than with intervening wake, although both associations were significant. This relationship, previously reported in rodents and small human samples, was reproduced across large clinical cohorts, including individuals with narcolepsy type 1. Inter-REM NREM duration was also bimodally distributed, with the 20–25 min valley closely preceding the first N3 epoch, suggesting a putative inhibition window.

While the homeostatic regulation of NREM-sleep pressure has been modelled with quantitative success [57, 58] and partially mapped onto identifiable molecular substrates [59], the homeostatic regulation of REM-sleep pressure (HREMSP) remains incompletely characterized [60]. Benington and Heller proposed that, over the course of a single night, HREMSP accumulates primarily during NREM sleep and dissipates during REM sleep [61, 62]. However, subsequent REM-deprivation studies support the idea that HREMSP also accumulates during wakefulness, motivating a distinction between short-term regulation of the NREM–REM cycle within a sleep period and long-term regulation of daily REM-sleep amount [63, 64]. Furthermore, REM-sleep deprivation induces REM rebound in suprachiasmatic-nucleus-lesioned animals lacking overt circadian rhythmicity in the absence of zeitgebers, supporting a homeostatic component independent of circadian gating [65].

Recent studies have reported that NREM duration within inter-REM intervals (IRIs) follows a bimodal distribution in both rodents [66] and humans [67], consistent with earlier sleep-interruption studies suggesting a temporal relationship between NREM duration and REM-sleep occurrence [68, 69]. Here we tested whether this bimodality is a robust feature of human sleep architecture and examined its temporal relationship to the emergence of N3 sleep. We then assessed whether subsequent REM-bout duration was more strongly associated with NREM duration within the IRI than with intervening wakefulness, and whether prior REM-bout duration influenced the subsequent IRI and the NREM sleep within it. These relationships were examined across large clinical cohorts, including individuals with narcolepsy type 1, a disorder caused by orexin deficiency in which REM propensity is greatly enhanced [70].




Figure 5: Inter-REM NREM-duration bimodality, N3 emergence, and REM-bout associations across cohorts. (a) Analysis schematic. 30-second hypnograms were generated by majority vote over 1-second hypnodensity epochs, and inter-REM intervals (IRIs) were extracted, including preceding REM bout (REM_{pre}), inter-REM interval comprising NREM time (N_s) and intervening wake (N_w), and subsequent REM bout (REM_{post}). (b) Distributions of $\log(N_s)$ across 30-min time-from-sleep-onset bins, showing changes in bimodality across the night. (c) Two-component Gaussian fit of $\log(N_s)$ and distribution of time to first N3. The N_s valley (23.1 min) precedes the median time to first N3 (26.5 min). (d) Slow-wave and spindle densities across $\log(N_s)$; the N_s valley coincides with the increasing slow-wave-density regime. (e) Covariate-adjusted ordinary least squares (OLS) coefficients relating REM_{pre} to subsequent N_s and IRI total duration, and N_s and N_w to REM_{post} across three cohorts. N_s was more strongly associated with REM_{post} than intervening wake was. (f) Consecutive short REM bouts lead to progressively shorter subsequent NREM durations. The orange curve shows the final IRI NREM duration after a long REM bout followed by $k = 1, 2, 3, 4$ short REM bouts (before a new long REM bout appears); the blue curve shows the corresponding control case, in which the long REM bout is followed by $k = 1, 2, 3, 4$ non-short REM bouts.

The primary analysis comprised 24,409 BioSerenity night recordings comprising 74,079 IRIs, of which 60,420 had complete demographic and AHI covariate data. Because segmentation choices affect bout and sleep-cycle definitions [71, 66], sensitivity analyses covered 81 REM/NREM segmentation configurations, were stratified by wake-after-sleep-onset (WASO) burden, and re-evaluated the primary findings for each successive sleep cycle. The analysis was extended to SSC (21,018 recordings, 69,502 IRIs, 25,652 with all demographic and AHI covariates) and the control subgroup of the CDH registry (1,181 recordings, 3,198 IRIs with all demographic and AHI covariates); the NT1 arm of CDH (501 recordings eligible for analysis, 2,433 IRIs with all demographic and AHI covariates) tested whether the same architecture held under disordered REM regulation. All covariate-adjusted regression models included age, sex, BMI, and AHI.

The distribution of $\log(N_s)$ (log duration of NREM in the IRI) within IRIs was bimodal: a two-Gaussian mixture fitted the data substantially better than a unimodal Gaussian (Figure 5b; difference in Bayesian information criterion, $\Delta\text{BIC} = -73,286$), a pattern that held across cohorts and across the 81 segmentation configurations. Its shape evolved across the night: IRIs with short NREM duration were improbable shortly after sleep onset, became progressively more probable with time asleep, and bimodality became less distinct towards morning (Figure 5b). The valley threshold of the distribution (23.1 min in BioSerenity) closely preceded the median duration at which the first N3 epoch appeared within the NREM bout (26.5 min), and coincided with the rising regime of slow-wave density. This was observed in all three cohorts (Figure 5c,d, Supplementary Figure 13a,c).

Across the three cohorts, we tested the previously reported association between the duration of a REM bout and the duration of the subsequent IRI NREM [62]. The log duration of the REM bout preceding an IRI (REM_{pre}) was associated with log total IRI duration. The association was stronger with $\log(N_s)$ (Figure 5e; covariate-adjusted OLS on the BioSerenity dataset with cluster-robust standard errors clustered at the recording level; CI derived from the cluster-robust covariance; REM_{pre} OLS coefficient for $\log_{10}(N_s)$ was 0.2477, SE 0.0064, $q < 0.001$). $\log(N_s)$ within an IRI was more strongly associated with $\log(\text{REM}_{\text{post}})$ than was log wake duration in the same interval, though both were significant (Figure 5e; covariate-adjusted, $\beta_{\log_{10}(N_s)} = 0.0738$, SE 0.0035, Shapley partial R^2 : 0.028, versus $\beta_{\log_{10}(N_{\text{wake}}+1)} = 0.0276$, SE 0.0026; Shapley partial R^2 : 0.011, both $q < 0.001$). Finally, following successive short REM bouts, subsequent NREM durations were progressively shorter than after a long REM bout followed by non-short bouts (Figure 5f; thresholds: long REM bout ≥ 10 min, short REM bout ≤ 5 min; Mann-Whitney U test, significant for $k = 1$ to 4). Sensitivity analyses showed that this finding was robust to the choice of short-bout threshold.

In the NT1 subgroup, the associations shown in Figure 5e and the across-night bimodality of $\log(N_s)$ were preserved, although the global valley threshold appeared earlier, at 18.7 min. Splitting patients by sleep-onset REM period (SOREMp+, first transition to REM below 15 min, aligned with American Academy of Sleep Medicine [AASM] recommendations, versus SOREMp-) did not reveal any difference in bimodality dynamics (Supplementary Figure 13b).

Three findings emerged consistently across cohorts and sensitivity configurations. First, IRI NREM duration was robustly bimodal, with a valley preserved across cohorts (23.1 min for BioSerenity, 23.3 min for the CDH control subgroup, and 21.4 min for SSC). Across the night, the IRI NREM duration distribution drifted towards more permissive short-bout regimes as sleep progressed.

Second, this valley coincided with the rising regime of slow-wave density and closely preceded the first N3 epoch. This temporal association is compatible with, but does not establish, a transient inhibition of the NREM-to-REM transition that would favour slow-wave consolidation. Because the pattern persisted into the late night when N3 is rare, it appears to be independent of the actual occurrence of N3. The observed pattern was also inconsistent with the proposed self-inhibition of REM sleep at its termination [72], as early NREM-to-REM transitions are more probable than transitions within the 20–25 min window throughout the night. In NT1 the same associations and bimodality held, with the valley and first N3 epoch appearing earlier than in controls.

Third, across several large clinical cohorts and in a population with disordered REM regulation, the reported stronger association of REM duration with NREM duration compared with wake during the IRI held [62]. Over a single night, the association of REM_{post} with IRI NREM duration was stronger than with IRI wake duration. Yet both were significant, which is compatible with the short- and long-term processes of HREMSP as suggested by Franken [64]. Moreover, because wake time (16 h) typically exceeds NREM time (6.5 h) over 24 h by almost threefold, the net contributions of wake and NREM could balance out to end up becoming equivalent over a full day. While the $\text{REM}_{\text{pre}}-N_s$ association is well established across species, this association between IRI NREM duration and subsequent REM-bout duration has not been reported in earlier rodent and human studies [61, 73, 74]. The effect we resolve is small, suggesting stronger coupling between REM_{pre}

and N_s than between N_s and REM_{post} , and may have been beyond the reach of much smaller cohorts. A direct test would require an independent marker, such as the rate of REM-onset attempts under selective deprivation [64].

The extent to which these associations inform HREMSP itself warrants caution, since our outcome is REM-bout duration rather than REM pressure. HREMSP has no established marker [66]; proposed measures are either timing-based [61, 75] or, as in our analysis, amount-based, where REM_{post} duration is treated as an index of accumulated propensity. This has been the working assumption under which probabilistic propensity measures have been validated in rodents [66, 73]. Within these models, the REM hourglass is proposed to govern both transition probability and episode consolidation [76].

The stronger β for $REM_{pre}-N_s$ than for N_s-REM_{post} could also suggest an alternative interpretation: that an NREM sleep debt accumulates during REM sleep and is subsequently discharged during NREM sleep. Moreover, circadian and homeostatic components cannot be fully separated within one night, despite the cycle-by-cycle analysis. Forced-desynchrony protocols [77] and more interventional models in suprachiasmatic-nucleus-lesioned animals would be needed to investigate this hypothesis further.

Expert note on collaboration

The operating hypothesis, that the bimodal distribution of inter-REM interval (IRI) durations reflects a robust architectural feature, emerged through human–AI collaboration. The investigator initially sought to extend an existing REM-homeostasis framework into a two-process-style model, while the agent’s iterations surfaced a weak fit for that model alongside a highly robust IRI bimodality, redirecting the investigation towards the latter. The investigator selected the three target cohorts (BioSerenity, SSC, and CDH), fixed the primary covariate set (age, sex, BMI, AHI), and pre-specified the sensitivity scope. Within these constraints, AI SLEEP CO-SCIENTIST designed and executed the bout-segmentation and statistical pipelines, performed the cross-cohort analyses and the NT1 application, generated all figures, and refined emerging sub-hypotheses across eight iterative rounds of investigator feedback. The agent was most responsive early in the process; broadly scoped instructions were occasionally misinterpreted or incompletely retained during later rounds. The principal benefit was compression of the exploratory phase: a relevant prompt, grounded in the field literature, could be prepared in less than an hour and dispatched overnight, with initial results available the following morning to determine whether the direction warranted further investigation.

Transient oscillation disruption in narcolepsy type 1

Main finding

A comprehensive scan of transient oscillations, defined as brief EEG bursts spanning 4–25 Hz, identified two robust, multi-site signatures: a frontal fast-sigma (spindle-band) deficit during N2/N3 sleep, independently validated using YASA spindle density; and a coherent redistribution towards slower rhythms, with elevated theta during Wake/NREM and low-alpha during Wake, N1, and REM. Beyond these robust signatures, slow-oscillation power was reduced across Wake, N1, N2, and REM, while the reorganization of transient-oscillation coupling to slow-oscillation phase was more site-sensitive and is therefore considered exploratory.

Conventional polysomnographic staging resolves sleep at the level of 30-second epochs, whereas the pathophysiology of narcolepsy type 1, which involves the selective loss of hypocretin (orexin) neurons, is expected to perturb sleep microstructure at much finer timescales. The biological question was whether a comprehensive scan of transient oscillations (TOs), brief, spectrally resolved EEG bursts detected using the Dynamic Oscillation Toolbox (DynamO) [33] across the full time–frequency plane (4–25 Hz), would reveal coherent biological structure beyond the handful of spindle measures examined in prior single-site work, and whether such structure could be recovered even when individual effects fell below the significance threshold.

Transient oscillations are brief, localized bursts of oscillatory activity in the sleep EEG that appear as peaks in the time–frequency representation of the signal. They reflect transient increases in rhythmic neural activity at a characteristic frequency before dissipating. Sleep spindles (12–16 Hz bursts lasting 0.5–2 s) are the best-known example, but TOs encompass a broader class of events spanning a wide range of frequencies (4–25 Hz), including theta (4–8 Hz), alpha (8–11 Hz), sigma (9–16 Hz), and beta (16–25 Hz). The key insight is




Figure 6: Transient-oscillation signatures of narcolepsy type 1. Throughout, head-maps in (d–g) are filled by effect size, Cohen's d (NT1 – control): *red* = higher in NT1, *blue* = lower in NT1, white ≈ 0 , on the diverging scale shown in Supplementary Figure 14 (–2 to +2); a filled dot marks FDR $q < 0.05$. In (h,i) colour instead encodes channel identity, not effect size. (a) Anatomy of a transient oscillation (TO) riding the up-state of a slow oscillation (SO), defining the per-TO features: *duration*, *prominence*, *SO-power*, *SO-phase*, and *rate* (TO count per stage); the inset contrasts a control spindle with the reduced NT1 spindle (SO 0.3–1.5 Hz, TO 4–25 Hz). The SO-phase dial sets the head-map convention. (b) Validation with an independent detector: Cohen's d (NT1 versus control) for fast- σ TO rate (solid) and YASA spindle density (hatched) at F4, C3, C4, O2. (c) Per-subject Pearson correlation between fast- σ TO rate and YASA spindle density. (d–g) Topographic detail for representative direction-based patterns: (d) P1 fast- σ duration reduction (N2); (e) P2 θ rate excess (Wake); (f) P3 low- α rate increase (REM); (g) P7 low- α SO-power (Wake). In (g), SO-power is reduced (blue) at all channels except F3, which shows a slight increase (red, $d = +0.28$); this left-frontal sparing is the defining feature of pattern P7. (h,i) N2 SO-phase angle (P9) for θ and slow- σ : each head-map arrow gives the mean SO-phase (coloured = NT1, grey = control; dashed when reliability is low, $\kappa < 0.7$), read against the dial in (a); arrow *colour* denotes the channel (F3, F4, C3, C4, O1, O2), not effect size. Panel titles are underlined in the colour of their pattern (P1–P9); the comprehensive feature-by-region scan from which these patterns are drawn, and the shared colour scale, are shown in Supplementary Figure 14.

that traditional spindle detectors apply amplitude and duration criteria historically calibrated to events visible by eye and therefore select only a high-amplitude subset of a broader continuum of transient events [78, 33].

The study was carried out end to end in the CDH cohort, a six-site clinical PSG archive comprising 482 individuals with NT1 and 417 controls [28, 29]. Six TO features (rate, prominence, duration, SO-phase coupling strength, SO-phase angle, and SO-power; Figure 6a) were computed for each channel, sleep stage, and frequency band and residualized for age, sex, BMI, and site. All 900 features were then tested with Benjamini–Hochberg FDR correction (Supplementary Figure 14). Two complementary strategies were applied: conventional FDR-based grouping and a direction-based coherent-pattern analysis that identifies feature-stage combinations with consistent effect directions across channels and frequency bands, even when individual effects are subthreshold. Expert feedback then guided the synthesis of these findings into five biologically interpretable themes (Figure 6).

Two themes emerged as robust NT1-associated signatures. First, NREM sleep showed a fast-sigma deficiency: spindle-band (12–16 Hz) TOs were less frequent and, with larger effect sizes, shorter and lower in amplitude during N2/N3, with the strongest effects observed frontally (F4 N2 rate: $d = -0.78$, 95% CI –1.03 to –0.50, $q = 4.8 \times 10^{-9}$; F4 N2 duration: $d = -1.23$, 95% CI –1.54 to –0.87, $q = 2.2 \times 10^{-15}$, the largest effect in the scan outside the circular SO-phase measures; $n = 102$ NT1/191 controls; Figure 6d). This deficiency was further validated using YASA spindle density (F4: $d = -0.49$, 95% CI –0.82 to –0.16, $n = 64/112$; C3: $d = -0.30$, 95% CI –0.48 to –0.11, $n = 223/284$; Figure 6b,c).

Second, oscillatory activity was redistributed. Theta rate and duration increased during Wake and NREM sleep, with the largest effect observed at F4 during Wake (rate $d = +1.07$, 95% CI $+0.77$ to $+1.35$, $q = 2.7 \times 10^{-12}$; Figure 6e). Low-alpha rate increased consistently across Wake, N1, and REM (18/18 features showed positive effects; mean $d = +0.35$; Figure 6f), whereas beta rate and duration decreased (F4 Wake rate: $d = -0.80$, 95% CI -1.06 to -0.56 , $q = 7.4 \times 10^{-9}$). Together, these findings indicate a coherent shift towards slower oscillatory modes, consistent with the state-boundary instability associated with hypocretin loss [79, 80].

The direction-based analysis assessed whether the sign of an effect was consistent across all six channels within a frequency band, regardless of whether individual features passed FDR correction. This analysis recovered coherent structure that significance thresholds alone missed, and that may warrant further investigation. The most consistent signal in the dataset was a cross-stage reduction in slow-oscillation (SO) power: every N1 and N2 feature showed the same effect direction (100% directional consistency), as did the large majority of Wake and REM features (Supplementary Figure 14), with left-frontal sparing during Wake and REM (Figure 6g). A second emergent theme was coupling reorganization. During N2, the preferred SO phase of every non-theta band shifted coherently, with the largest effects observed centrally (low-alpha $d = 1.42$ – 1.49 ; Figure 6h,i), whereas theta alone broke this coherence. One possible interpretation, suggested by AI SLEEP CO-SCIENTIST but not tested here, is that theta activity arises from hippocampal–cortical rather than thalamocortical generators.

Two additional patterns were considered exploratory. The first was a weak prominence asymmetry (P6; Supplementary Figure 14): 129 of 150 prominence features showed a consistent left/right split, with greater prominence in NT1 than in controls at F3, C3, and O2 and lower prominence at F4, C4, and O1. The second comprised large N2 SO-phase-angle shifts (P9; Figure 6h,i), which attenuated after stratification by site. Neither pattern is interpreted as an established finding.

Taken together, centrofrontal-channel fast-sigma reduction and theta excess, mostly detectable in N2, emerge as robust, multi-site NT1-associated signatures of sleep microstructure. The physiological significance of this shift from sigma to theta is uncertain; most PSG studies to date have focused on N1, Wake, and REM sleep, where such abnormalities are typically found [28]. In contrast, spindle density has generally been reported as normal in NT1 [81]. The finding here suggests that abnormal distribution features beyond spindle density are present in NT1. Spindles are generally considered an electrophysiological biomarker of cognitive capacity and thalamocortical circuit integrity. Such a defect could thus partially explain cognitive difficulties beyond sleepiness that are reported by patients with NT1/orexin deficiency [82]. Both fast-sigma reduction and theta excess sit within a broader redistribution of oscillatory rate and diminished slow-oscillation background, while the SO-phase-coupling findings, which attenuate under site stratification, are reported as hypothesis-generating. This case illustrates the value of three elements in combination: a comprehensive scan, a direction-based analysis that treats consistent effect sign as evidence, and expert synthesis that both organizes the findings and tests them against confounds. It also revealed a new microarchitectural abnormality in NT1.

Expert note on collaboration

The investigator set the broad direction of a sub-epoch TO scan in the CDH cohort, and the agent proposed analytical strategies to support it. The investigator then fixed the design: the NT1-versus-control contrast, the six TO feature families (rate, prominence, duration, SO-phase coupling strength, SO-phase angle, SO-power), the covariate set (age, sex, BMI, site), and a dual FDR-corrected and direction-based analysis. Within that design, AI SLEEP CO-SCIENTIST computed all features, ran both analyses, screened artefacts, performed the site-stratified confound checks, and generated the figures. The study was run as several parallel branches with deliberately different methodological choices, each refined over roughly five to eight rounds of feedback; only findings that reproduced across branches were reported. Feedback was mostly methodological: sensitivity analyses (e.g. site-stratification), prominence artefact removal (16 outliers removed), and scoping or figure edits. It was especially useful for probing AI SLEEP CO-SCIENTIST’s non-obvious choices. The earliest iterations needed the most attention; once the analytical contract stabilized, changes were usually integrated cleanly, and on long chains a fresh restart proved more effective than continued steering.

Discussion

We applied AI SLEEP CO-SCIENTIST to five case studies in four clinical and epidemiological PSG cohorts, spanning disease risk, clinical phenotyping, and mechanism. Each case study recovered established features of

sleep physiology before addressing a question that had not previously been examined at this scale: coupling between organ systems during sleep, the domain composition of physiological ageing, the temporal organization of cortical arousals, NREM–REM cycling structure, and the organization of transient oscillations in narcolepsy type 1. All five were refined through successive rounds of expert feedback, and selected findings were assessed through independent reruns, external validation, or cross-site replication where the available data and study design allowed. The evaluation therefore reflects repeated, expert-reviewed workflows rather than the output of a single unattended run.

In both disease-risk studies the disease signal was specific rather than global. In the coupling study, alterations were specific to disorder, sleep stage, and subnetwork: reduced N2 brain–brain coupling was associated with a higher risk of subsequent Parkinson’s disease and Alzheimer’s disease diagnoses, and reduced brain–heart coupling with subsequent Parkinson’s disease, Alzheimer’s disease, and dementia. Both are biologically plausible. Progressive loss of cardiac noradrenergic sympathetic innervation is characteristic of Parkinson’s disease and can occur independently of nigrostriatal dopaminergic degeneration [38]. One possible explanation for the brain–brain findings is that deficits in dopamine in Parkinson’s disease and acetylcholine in Alzheimer’s disease, together with proteinopathy-related synaptic and interneuron impairment, disrupt the excitatory–inhibitory balance and structural connectivity required for coordinated oscillatory coupling. Cross-site concordance of these effect-size patterns was only moderate. In the sleep-ageing study the specificity appeared at the level of organ systems instead. Decomposing PSG ageing into physiologically defined domains yielded a more interpretable representation that generalized better across cohorts than an unconstrained model. This suggests that future sleep-ageing models should preserve modular structure across neural, respiratory, autonomic, and motor systems, and should test adaptive or conditionally weighted fusion rather than simply increasing feature dimensionality. The resulting residual should not be interpreted as a universal biological-age axis: depending on the outcome, it may reflect shared age-related vulnerability, disease-specific pathophysiology, consequences of established disease, or combinations of these processes [43]. Disease-tuned scores remain useful secondary risk models, but in this analysis they did not replace the age-derived late-fusion residual as the primary ageing construct.

Our findings refine current models of COMISA by suggesting that its distinguishing physiology lies less in an exaggerated burden of respiratory events than in an altered response to those events [83, 55, 49]. The prolonged wakefulness and greater irregularity following cortical arousals support the hypothesis that impaired re-establishment of stable sleep, rather than increased respiratory instability itself, may be a defining feature of the disorder [55, 49]. This interpretation is consistent with concepts of insomnia-related hyperarousal while also explaining why COMISA remained closely aligned with OSA across most sleep-microstructure measures [84, 55]. More broadly, these results highlight the value of moving beyond conventional PSG summary metrics towards dynamic measures of cortical arousal organization, which may better capture clinically relevant heterogeneity in sleep disorders [27, 85]. Although the observational design precludes causal inference and the degree of OSA skew varied across cohorts, the consistent prominence of post-arousal dynamics suggests these measures warrant investigation as markers for stratifying COMISA and for identifying patients who might benefit from interventions targeting both respiratory events and post-arousal sleep stabilization.

Within inter-REM intervals, subsequent REM-bout duration tracked intervening NREM duration more closely than intervening wake, reproducing a relationship previously reported in rodents and small human samples and holding under the disordered REM regulation of narcolepsy type 1. That both associations were significant is compatible with separate short-term and long-term components of REM homeostasis [64], and because waking time over 24 hours exceeds NREM time roughly threefold, their net daily contributions could still be comparable. Inter-REM NREM duration was also bimodally distributed, with the valley closely preceding the first N3 epoch and coinciding with rising slow-wave density; this temporal association is compatible with, but does not establish, a transient inhibition of the NREM-to-REM transition that would favour slow-wave consolidation. The extent to which these associations inform REM homeostasis itself warrants caution. The outcome is REM-bout duration rather than REM pressure, for which no established marker exists [66]; circadian and homeostatic contributions cannot be separated within a single night; and an alternative reading, in which NREM debt accumulates during REM sleep and is discharged during NREM sleep, is also compatible with the observed asymmetry. Forced-desynchrony protocols [77] or interventional animal models would be needed to distinguish them. In narcolepsy type 1, the fast-sigma deficit and centrofrontal theta excess point to microarchitectural pathology that 30-second staging does not resolve. Spindle density has generally been reported as normal in this disorder [81], so the abnormality involves properties of spindle-band events beyond their density. Because spindles are considered a marker of cognitive capacity and thalamocortical circuit integrity, such a defect could partially explain the cognitive difficulties patients report beyond sleepiness [82].

Obtaining these results depended on an expert-directed division of labour that shifted across the pipeline. At the hypothesis stage the most productive directions arose when experts seeded the system with relevant literature and scientific intuition, after which AI SLEEP CO-SCIENTIST expanded the candidate space and assessed which questions the data could support. A focused, reference-seeded prompt could be dispatched overnight, letting an investigator survey the conceptual neighbourhood of an idea before committing to it. During preprocessing, pipelines were developed and tested on sample recordings, with expert review before full-cohort execution aligning methodological choices with domain conventions and checking the physiological plausibility of the resulting features. Execution was the most autonomous stage: once the analytical contract was fixed, the study proceeded through staged analysis, internal critique, and reporting, with the expert intervening at checkpoints rather than directing every step. Two design choices made this workable at scale. First, because the cohorts hold more than 50TB of raw signal, expensive preprocessing is separated from hypothesis-specific analysis and validated features are cached for reuse, amortizing the cost of raw-signal processing across studies. Second, because PSG spans systems with distinct noise characteristics and clinical conventions, cohort definitions, canonical splits, channel conventions, and analytical guidance are encoded in the shared substrate rather than restated in every prompt, which narrows the space of plausible but inappropriate analytical choices.

Several lessons emerged across the investigators who used the system, and each of them changed how it was built or used over the course of the case studies. Prompt structure mattered: broad or multi-part requests produced diffuse analyses, whereas narrow prompts with a single objective yielded more focused and reproducible outputs, and investigators therefore converged on decomposing complex aims into self-contained tasks. Long-horizon work degraded in a characteristic way, with the system drifting from the original question, carrying forward early assumptions, or letting an error introduced in one iteration shape later analyses as earlier constraints were revised or lost. Restarting a step with a consolidated specification proved more effective than appending further feedback, and became the standard response once an analytical contract had drifted. Neither transparency nor fidelity to the evidence was automatic. Early runs had to be asked explicitly for the cohort-flow tables, exclusion summaries, and intermediate artefacts needed to evaluate a result, and successive rounds of feedback and revision of the system made these outputs a routine part of every run. The provenance artefacts described in the Methods, including the cohort-flow record whose arithmetic is checked mechanically, address the same gap. Contributions also ran in both directions: within an expert-defined contract the system proposed useful methodological changes, including the age-bias correction adopted in the sleep-ageing study, functioning less as an execution tool than as a collaborator whose suggestions could be evaluated and revised.

There are several limitations in this work. The analyses are retrospective and rest on clinical-referral or epidemiological cohorts rather than general-population samples, making them subject to the ascertainment biases inherent in how these recordings were collected. The reported associations should be treated as hypothesis-generating rather than causal, and require designs that emulate a target trial, and ultimately prospective or interventional studies, for confirmation. They also rest on single-night, in-laboratory PSG, which is subject to the first-night effect and does not capture night-to-night variability at home; a single recording is an imperfect estimate of habitual sleep, and longitudinal designs are needed to separate state from trait [86]. The case studies vary in their proximity to clinical translation: the COMISA arousal-recovery phenotype and the sleep-age residuals rely on routinely derived PSG features, whereas network-physiology coupling and transient-oscillation analysis require full-bandwidth multichannel recordings and specialized post-processing that most clinical systems do not support. None has been evaluated as a clinical decision tool, and actionability will require prospective studies showing that acting on these measures improves outcomes. The scope is bounded in other ways: the hypothesis and execution stages operated mainly on precomputed features rather than raw signals, and five case studies sample only a small region of a vast hypothesis space. Finally, automated critique can catch procedural and internal-consistency errors but is not a dependable arbiter of scientific novelty or importance, so judging whether an open-ended finding is genuinely new still rests with domain experts. The safeguards reduce but do not eliminate hallucination, and subtle fabricated or misattributed content can survive automated checks, particularly as errors accumulate across revisions. Behaviour also varied across model versions and capability tiers, which presents a reproducibility challenge as deployed models change.

Several directions could extend this work. Larger and more representative cohorts, including general-population and prospective cohorts with multi-night recordings, would help separate trait from state, and samples collected under protocols that dissociate homeostatic from circadian influences, such as constant-routine and forced-desynchrony designs, would allow the mechanistic questions raised here to be tested directly. Additional patient cohorts are also needed, as some pathologies uniquely challenge existing assumptions.

A next generation of the system could complement its curated feature library with direct access to raw signals and learned representations, including SleepFM embeddings [8], enabling patterns not captured by predefined features to be discovered and interrogated. Ultimately, cross-modal retrieval across PSG, electronic health records, imaging, wearables, and genomics could support hypotheses spanning multiple levels of physiology. The deeper challenge is to widen the portion of the scientific process that can be automated without sacrificing rigour, where the bottleneck is evaluation rather than generation. Open problems include a rigorous automated judge of methodological soundness and claim calibration, a validator that independently reproduces and stress-tests a result rather than confirming it as written, and reliable measures of genuine hypothesis novelty. One resource for that effort may be a by-product of the workflow itself: because every run preserves the collaboration trajectory behind a study, including the questions posed, the corrections made, and the judgements that shaped the analysis, these records could in principle train future models to act as more effective scientific collaborators.

Acknowledgements

The Human Sleep Project is supported by a grant from the National Heart, Lung, and Blood Institute of the NIH (R01HL161253).

R.T. is supported by Knight-Hennessy Scholars funding. E.M. is supported by a grant from the National Heart, Lung, and Blood Institute of the NIH (R01HL161253). J.Z. is supported by funding from the Chan-Zuckerberg Biohub. We thank Topi Niemi and Esa Räsänen for their help with the preprocessing steps to calculate the HRV features.

Author Contributions Statement

R.T., E.M., and J.Z. conceived the project. R.T. led the study, developed the agent system, and coordinated the five case studies. R.T. and M.R.K. built the initial iteration of the agent infrastructure. U.H., R.G., A.S., A.B.-K., and E.H. preprocessed the expert-led datasets. R.G. led the REM-sleep regulation case study, A.S. led the transient-oscillation study, A.B.-K. led the sleep-ageing study, U.H. led the COMISA study, and M.S. led the network-physiology study. H.G.Z. provided clinical expertise and guidance on agentic system design. E.C.L. provided domain-expert feedback that guided the iterative refinement of AI-generated hypotheses, analyses, and scientific interpretation. E.M. and J.Z. supervised the project. All authors contributed to writing and copyediting the manuscript.

References

- [1] Francesco P Cappuccio, Lanfranco D’Elia, Pasquale Strazzullo, and Michelle A Miller. Sleep duration and all-cause mortality: a systematic review and meta-analysis of prospective studies. *Sleep*, 33(5):585–592, 2010.
- [2] Julian Lim and David F Dinges. A meta-analysis of the impact of short-term sleep deprivation on cognitive variables. *Psychological bulletin*, 136(3):375, 2010.
- [3] Chiara Cirelli and Giulio Tononi. Is sleep essential? *PLoS biology*, 6(8):e216, 2008.
- [4] Allan Rechtschaffen, Marcia A Gilliland, Bernard M Bergmann, and Jacqueline B Winter. Physiological correlates of prolonged sleep deprivation in rats. *Science*, 221(4606):182–184, 1983.
- [5] Matthew Walker. *Why we sleep: Unlocking the power of sleep and dreams*. Simon and Schuster, 2017.
- [6] Richard B Berry, Rita Brooks, Charlene E Gamaldo, Susan M Harding, Carole Marcus, Bradley V Vaughn, et al. The AASM manual for the scoring of sleep and associated events. *Rules, Terminology and Technical Specifications, Darien, Illinois, American Academy of Sleep Medicine*, 176(2012):7, 2012.
- [7] Samaneh Nasiri, Wolfgang Ganglberger, Thijs Nassi, Erik-Jan Meulenbrugge, Valdery Moura Junior, Manohar Ghanta, Aditya Gupta, Katie L Stone, Magnus Ruud Kjaer, Oliver Sum-Ping, et al. CAISR: achieving human-level performance in automated sleep analysis across all clinical sleep metrics. *Sleep*, 48(8):zsaf134, 2025.
- [8] Rahul Thapa, Bryan He, Magnus Ruud Kjaer, Hyatt Moore, Gauri Ganjoo, Emmanuel Mignot, and James Zou. SleepFM: multi-modal representation learning for sleep across brain activity, ECG and respiratory signals. *arXiv preprint arXiv:2405.17766*, 2024.
- [9] Rahul Thapa, Magnus Ruud Kjaer, Bryan He, Ian Covert, Hyatt Moore IV, Umaer Hanif, Gauri Ganjoo, M Brandon Westover, Poul Jennum, Andreas Brink-Kjaer, et al. A multimodal sleep foundation model for disease prediction. *Nature Medicine*, pages 1–11, 2026.
- [10] Magnus Ruud Kjaer, Rahul Thapa, Gauri Ganjoo, Hyatt Moore IV, Poul Joergen Jennum, Brandon M Westover, James Zou, Emmanuel Mignot, Bryan He, and Andreas Brink-Kjaer. Stanford Sleep Bench: Evaluating polysomnography pre-training methods for sleep foundation models. *arXiv preprint arXiv:2512.09591*, 2025.
- [11] Zongzhe Xu, Zitao Shuai, Eideen Mozaffari, Ravi S Aysola, Rajesh Kumar, and Yuzhe Yang. SleepLM: Natural-language intelligence for human sleep. *arXiv preprint arXiv:2602.23605*, 2026.
- [12] Daniil A Boiko, Robert MacKnight, Ben Kline, and Gabe Gomes. Autonomous chemical research with large language models. *Nature*, 624(7992):570–578, 2023.
- [13] Chris Lu, Cong Lu, Robert Tjarko Lange, Yutaro Yamada, Shengran Hu, Jakob Foerster, David Ha, and Jeff Clune. Towards end-to-end automation of AI research. *Nature*, 651(8107):914–919, 2026.
- [14] Ludovico Mitchener, Angela Yiu, Benjamin Chang, Mathieu Bourdenx, Tyler Nadolski, Arvis Sulovari, Eric C Landsness, Daniel L Barabasi, Siddharth Narayanan, Nicky Evans, et al. Kosmos: An AI scientist for autonomous discovery. *arXiv preprint arXiv:2511.02824*, 2025.
- [15] Kexin Huang, Serena Zhang, Hanchen Wang, Yuanhao Qu, Yingzhou Lu, Ryan Li, Yusuf Roohani, Lin Qiu, Shiyi Cao, Gavin Li, et al. Autonomous biomedical research with an artificial intelligence agent. *Science*, page eadz4351, 2026.
- [16] Jiabin Tang, Lianghao Xia, Zhonghang Li, and Chao Huang. Ai-researcher: Autonomous scientific innovation. *Advances in Neural Information Processing Systems*, 38:9481–9520, 2026.
- [17] Pan Lu, Yifan Gao, William G Peng, Haoxiang Zhang, Kunlun Zhu, Elektra K Robinson, Qixin Xu, Masakazu Kotaka, Harrison G Zhang, Bingxuan Li, et al. Eubiota: Modular agentic AI for autonomous discovery in the gut microbiome. *bioRxiv*, pages 2026–02, 2026.
- [18] Harrison G Zhang, Peter Eckmann, Jiacheng Miao, Andrew B Mahon, and James Zou. The virtual biotech: A multi-agent AI framework for therapeutic discovery and development. *bioRxiv*, pages 2026–02, 2026.
- [19] Alexander Novikov, Ng  n V  , Marvin Eisenberger, Emilien Dupont, Po-Sen Huang, Adam Zsolt Wagner, Sergey Shirobokov, Borislav Kozlovskii, Francisco JR Ruiz, Abbas Mehrabian, et al. Alphaevolve: A coding agent for scientific and algorithmic discovery. *arXiv preprint arXiv:2506.13131*, 2025.
- [20] Juraj Gottweis, Wei-Hung Weng, Alexander Daryin, Tao Tu, Petar Sirkovic, Artiom Myaskovsky, Grzegorz Glowaty, Felix Weissenberger, Alessio Orlandi, Dan Popovici, et al. Accelerating scientific discovery with Co-Scientist. *Nature*, pages 1–3, 2026.

- [21] Shanghua Gao, Ada Fang, Yepeng Huang, Valentina Giunchiglia, Ayush Noori, Jonathan Richard Schwarz, Yasha Ektefaie, Jovana Kondic, and Marinka Zitnik. Empowering biomedical discovery with AI agents. *Cell*, 187(22):6125–6151, 2024.
- [22] Tal Ifargan, Lukas Hafner, Maor Kern, Ori Alcalay, and Roy Kishony. Autonomous LLM-driven research—from data to human-verifiable research papers. *NEJM AI*, 2(1):A10a2400555, 2025.
- [23] Kyle Swanson, Wesley Wu, Nash L Bulaong, John E Pak, and James Zou. The Virtual Lab of AI agents designs new SARS-CoV-2 nanobodies. *Nature*, 646(8085):716–723, 2025.
- [24] Qichen Li, Shenghan Wen, Haoqi Sun, Wolfgang Ganglberger, Ayush Tripathi, Niels Turley, Samuel Waters, Arnav Gupta, Aditya Gupta, Manohar Ghanta, Bruce Nearing, Han Wu, Katie L. Stone, Chad Robichaux, Zhiyong Zhang, Qiao Li, Gauri Ganjoo, Christine Tsien Silvers, Bharath Gunapati, Kiran Maski, Samaneh Nasiri, Dennis Hwang, Lynn Marie Trotti, Umakanth Katwa, Gari D. Clifford, Emmanuel Mignot, Robert J. Thomas, and M. Brandon Westover. The Human Sleep Project (version 3.0). <https://doi.org/10.60508/m3sw-rz13>, 2026. Version 3.0.
- [25] Umaer Hanif, Ulysse Gimenez, Alyssa Cairns, Daniel Lewin, Nisar Ashraf, and Emmanuel Mignot. Automatic detection of chronic insomnia from polysomnographic and clinical variables using machine learning. In *2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 1–5. IEEE, 2023.
- [26] Umaer Hanif, Alyssa Cairns, Vincent Mysliwiec, Ruggero G Bettinardi, Maxime Elbaz, Ulysse Gimenez, and Emmanuel JM Mignot. Associations between self-reported parasomnias and psychiatric illness in 370,000 patients with sleep disorders. *Psychiatry and Clinical Neurosciences*, 78(11):667–677, 2024.
- [27] Umaer Hanif, Veronica Guadagni, Kari Lambing, Ulysse Gimenez, Poul Jennum, Alexander Sweetman, and Vincent Mysliwiec. The use of the odds ratio product and self-reported data to detect comorbid insomnia and sleep apnea. *SLEEPJ*, 49(2):zsaf198, 2026.
- [28] Jens B. Stephansen, Alexander N. Olesen, Mads Olsen, Aditya Ambati, Eileen B. Leary, Hyatt E. Moore, Oscar Carrillo, Ling Lin, Fang Han, Han Yan, Yun L. Sun, Yves Dauvilliers, Sabine Scholz, Lucie Barateau, Birgit Hogl, Ambra Stefani, Seung Chul Hong, Tae Won Kim, Fabio Pizza, Giuseppe Plazzi, Stefano Vandi, Elena Antelmi, Dimitri Perrin, Samuel T. Kuna, Paula K. Schweitzer, Clete Kushida, Paul E. Peppard, Helge B. D. Sorensen, Poul Jennum, and Emmanuel Mignot. Neural network analysis of sleep stages enables efficient diagnosis of narcolepsy. *Nature Communications*, 9(1):5229, 2018.
- [29] Guo-Qiang Zhang, Licong Cui, Remo Mueller, Shiqiang Tao, Matthew Kim, Michael Rueschman, Sara Mariani, Daniel Mobley, and Susan Redline. The National Sleep Research Resource: towards a sleep data commons. *Journal of the American Medical Informatics Association*, 25(10):1351–1358, 2018.
- [30] Mathias Perslev, Sune Darkner, Lykke Kempfner, Miki Nikolic, Poul Jørgen Jennum, and Christian Igel. U-Sleep: resilient high-frequency sleep staging. *NPJ digital medicine*, 4(1):72, 2021.
- [31] Andreas Brink-Kjaer, Alexander Neergaard Olesen, Paul E Peppard, Katie L Stone, Poul Jennum, Emmanuel Mignot, and Helge BD Sorensen. Automatic detection of cortical arousals in sleep and their contribution to daytime sleepiness. *Clinical Neurophysiology*, 131(6):1187–1203, 2020.
- [32] Magnus Ruud Kjaer, Umaer Hanif, Andreas Brink-Kjaer, Mads Olsen, Oliver Sum-Ping, Oscar Carrillo, Scott A Sands, Susan Redline, Katie L Stone, Poul Jennum, et al. Expert-level probabilistic breathing event detector informs phenotyping of sleep apnea. *Nature Communications*, 2026.
- [33] Patrick A Stokes, Preetish Rath, Thomas Possidente, Mingjian He, Shaun Purcell, Dara S Manoach, Robert Stickgold, and Michael J Prerau. Transient oscillation dynamics during sleep provide a robust basis for electroencephalographic phenotyping and biomarker identification. *Sleep*, 46(1):zsac223, 2023.
- [34] M. Saibene, M. Ballegaard, Y. Gu, et al. A continuous network physiology analysis of brain–heart interactions in epileptic seizures. *Scientific Reports*, July 2026. Published 10 July 2026.
- [35] Raphael Vallat and Matthew P Walker. An open-source, high-performance tool for automated sleep staging. *eLife*, 10:e70092, 2021.
- [36] Amir Bashan, Ronny P Bartsch, Jan W Kantelhardt, Shlomo Havlin, and Plamen Ch Ivanov. Network physiology reveals relations between network topology and physiological function. *Nature communications*, 3(1):702, 2012.
- [37] Teemu Pukkila, Topi Niemi, and Esa Räsänen. Chronic heart failure detection based on long-term RR interval dynamics. *Journal of Electrocardiology*, 97:154242, 2026.

- [38] Diego Candia-Rivera, Marie Vidailhet, Mario Chavez, and Fabrizio De Vico Fallani. A framework for quantifying the coupling between brain connectivity and heartbeat dynamics: Insights into the disrupted network physiology in Parkinson’s disease. *Human Brain Mapping*, 45(5):e26668, 2024.
- [39] Andreas Brink-Kjaer, Eileen B Leary, Haoqi Sun, M Brandon Westover, Katie L Stone, Paul E Peppard, Nancy E Lane, Peggy M Cawthon, Susan Redline, Poul Jennum, et al. Age estimation from sleep studies using deep learning predicts life expectancy. *NPJ digital medicine*, 5(1):103, 2022.
- [40] Haoqi Sun, Luis Paixao, Jefferson T Oliva, Balaji Goparaju, Diego Z Carvalho, Kicky G van Leeuwen, Oluwaseun Akeju, Robert J Thomas, Sydney S Cash, Matt T Bianchi, et al. Brain age from the electroencephalogram of sleep. *Neurobiology of aging*, 74:112–120, 2019.
- [41] Haoqi Sun, Sasha Milton, Yi Fang, Hash Brown Taha, Shreya Shiju, Robert J Thomas, Wolfgang Ganglberger, Matthew P Pase, Timothy Hughes, Shaun Purcell, et al. Machine learning–based sleep electroencephalographic brain age index and dementia risk: an individual participant data meta-analysis. *JAMA Network Open*, 9(3):e261521, 2026.
- [42] Noor Adra, Lisa W Dümmer, Luis Paixao, Ryan A Tesh, Haoqi Sun, Wolfgang Ganglberger, Mike Westmeijer, Madalena Da Silva Cardoso, Anagha Kumar, Elissa Ye, et al. Decoding information about cognitive health from the brainwaves of sleep. *Scientific Reports*, 13(1):11448, 2023.
- [43] Carlos López-Otín, Maria A Blasco, Linda Partridge, Manuel Serrano, and Guido Kroemer. Hallmarks of aging: An expanding universe. *Cell*, 186:243–278, 1 2023.
- [44] *International classification of sleep disorders*. American Academy of Sleep Medicine, Darien, IL, 3 edition, 2014.
- [45] Alexander M Sweetman, Leon C Lack, Peter G Catcheside, Nick A Antic, Ching Li Chai-Coetzer, Simon S Smith, James A Douglas, and R Doug McEvoy. Developing a successful treatment for co-morbid insomnia and sleep apnoea. *Sleep medicine reviews*, 33:28–38, 2017.
- [46] Alexander Sweetman, Seamas Farrell, Douglas M Wallace, and Megan Crawford. The effect of cognitive behavioural therapy for insomnia in people with comorbid insomnia and sleep apnoea: A systematic review and meta-analysis. *Journal of Sleep Research*, 32(6):e13847, 2023.
- [47] Bastien Lechat, Sarah Appleton, Yohannes Adama Melaku, Kristy Hansen, R Doug McEvoy, Robert Adams, Peter Catcheside, Leon Lack, Danny J Eckert, and Alexander Sweetman. The association of co-morbid insomnia and sleep apnea with prevalent cardiovascular disease and incident cardiovascular events. *Journal of sleep research*, 31(5):e13563, 2022.
- [48] Alexander Sweetman, Leon Lack, R Doug McEvoy, Simon Smith, Danny J Eckert, Amal Osman, Jayne C Carberry, Douglas Wallace, Phuc D Nguyen, and Peter Catcheside. Bi-directional relationships between co-morbid insomnia and sleep apnea (COMISA). *Sleep medicine reviews*, 60:101519, 2021.
- [49] Elliot J Brooker, Shane A Landry, Luke DJ Thomson, Garun S Hamilton, Pedro R Genta, Sean PA Drummond, and Bradley A Edwards. Obstructive sleep apnea is a distinct physiological endotype in individuals with comorbid insomnia and sleep apnea. *Annals of the American Thoracic Society*, 20(10):1508–1515, 2023.
- [50] Marcela Yanagimori, Mariana D Fernandes, Michelle L Garcia, Paula G Scudeller, Carlos RR Carvalho, Bradley Edwards, Geraldo Lorenzi-Filho, and Pedro R Genta. Respiratory arousal threshold among patients with isolated sleep apnea and with comorbid insomnia (COMISA). *Scientific reports*, 13(1):7638, 2023.
- [51] Ervin Cotrik, Janete Hernandez, Viviane Castro, and Edilson Zancanella. Physiopathological correlations of comorbid insomnia and sleep apnoea (COMISA)—a systematic review and meta-analysis. *Sleep and Breathing*, 30(2):101, 2026.
- [52] Massimiliano de Zambotti, John Trinder, Alessandro Silvani, Ian M Colrain, and Fiona C Baker. Dynamic coupling between the central and autonomic nervous systems during sleep: a review. *Neuroscience & Biobehavioral Reviews*, 90:84–103, 2018.
- [53] Chiara Baglioni, Wolfram Regen, Armand Teghen, Kai Spiegelhalder, Bernd Feige, Christoph Nissen, and Dieter Riemann. Sleep changes in the disorder of insomnia: a meta-analysis of polysomnographic studies. *Sleep medicine reviews*, 18(3):195–213, 2014.
- [54] Aurélie M Stephan and Francesca Siclari. Reconsidering sleep perception in insomnia: from misperception to mismeasurement. *Journal of sleep research*, 32(6):e14028, 2023.

- [55] Bernice M Wulterkens, Lieke WA Hermans, Pedro Fonseca, Hennie CJP Janssen, Pauline V van Hirtum, Sebastiaan Overeem, and Merel M van Gilst. Heart rate response to cortical arousals in patients with isolated obstructive sleep apnea and with comorbid insomnia (COMISA). *Sleep and Breathing*, 28(2):735–744, 2024.
- [56] Yanyuan Dai, Yun Li, Elisabeth Heremans, Ulysse Gimenez, Umaer Hanif, and Emmanuel Mignot. Interactions between insomnia and obstructive sleep apnea. *medRxiv*, pages 2026–07, 2026.
- [57] Alexander A Borb and Peter Achermann. Sleep homeostasis and models of sleep regulation. *Journal of biological rhythms*, 14(6):559–570, 1999.
- [58] Pooja Rajdev, David Thorsley, Srinivasan Rajaraman, Tracy L Rupp, Nancy J Wesensten, Thomas J Balkin, and Jaques Reifman. A unified mathematical model to quantify performance impairment for both chronic sleep restriction and total sleep deprivation. *Journal of theoretical biology*, 331:66–77, 2013.
- [59] Valérie Mongrain, Marcos G Frank, and Tanya Leduc. Revisiting brain gene expression changes and protein modifications tracking homeostatic sleep pressure. *npj Biological Timing and Sleep*, 2(1):30, 2025.
- [60] Sung-Ho Park and Franz Weber. Neural and homeostatic regulation of REM sleep. *Frontiers in psychology*, 11:1662, 2020.
- [61] JOEL H Benington and H Craig Heller. REM-sleep timing is controlled homeostatically by accumulation of REM-sleep propensity in non-REM sleep. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 266(6):R1992–R2000, 1994.
- [62] Joel H Benington and H Craig Heller. Does the function of REM sleep concern non-REM sleep or waking? *Progress in neurobiology*, 44(5):433–449, 1994.
- [63] Adrián Ocampo-Garcés, Enrique Molina, Alberto Rodríguez, and Ennio A Vivaldi. Homeostasis of REM sleep after total and selective sleep deprivation in the rat. *Journal of Neurophysiology*, 84(5):2699–2702, 2000.
- [64] Paul Franken. Long-term vs. short-term processes regulating REM sleep. *Journal of sleep research*, 11(1):17–28, 2002.
- [65] Sarah W Wurts and Dale M Edgar. Circadian and homeostatic control of rapid eye movement (REM) sleep: promotion of REM tendency by the suprachiasmatic nucleus. *Journal of Neuroscience*, 20(11):4300–4310, 2000.
- [66] Alexander G Ginsberg, Madelyn Esther C Cruz, Franz Weber, Victoria Booth, and Cecilia G Diniz Behn. A predictive propensity measure to enter REM sleep. *Frontiers in Neuroscience*, 18:1431407, 2024.
- [67] Naghmeh Akhavan, Alexander G Ginsberg, Madelyn EC Cruz, Yunxi Yan, Shelby R Stowe, Dinesh Pal, Franz Weber, Cecilia G Diniz Behn, and Victoria Booth. A data-driven measure of REM sleep propensity for human and rodent sleep. *arXiv preprint arXiv:2604.01252*, 2026.
- [68] Tomoka Takeuchi, Kazuhiko Fukuda, and Timothy I Murphy. Elicitation of sleep-onset REM periods in normal individuals using the sleep interruption technique (SIT). *Sleep Medicine*, 3(6):479–488, 2002.
- [69] Y Sasaki, A Miyasita, T Takeuchi, M Inugami, K Fukuda, and K Ishihara. Effects of sleep interruption on body temperature in human subjects. *Sleep*, 16(5):478–483, 1993.
- [70] Emmanuel Mignot, Jamie Zeitzer, Fabio Pizza, and Giuseppe Plazzi. Sleep problems in narcolepsy and the role of hypocretin/orexin deficiency. In *Frontiers of Neurology and Neuroscience*. S. Karger AG, 2021.
- [71] Christine Blume and Christian Cajochen. ‘SleepCycles’ package for R – a free software tool for the detection of sleep cycles from sleep staging. *MethodsX*, 8:101318, 2021.
- [72] Olivier Le Bon. An asymmetrical hypothesis for the NREM-REM sleep alternation—what is the NREM-REM cycle? *Frontiers in Neuroscience*, 15:627193, 2021.
- [73] Sung-Ho Park, Justin Baik, Jiso Hong, Hanna Antila, Benjamin Kurland, Shinjae Chung, and Franz Weber. A probabilistic model for the ultradian timing of REM sleep in mice. *PLoS computational biology*, 17(8):e1009316, 2021.
- [74] Giuseppe Barbato and Thomas A Wehr. Homeostatic regulation of REM sleep in humans during extended sleep. *Sleep*, 21(3):267–276, 1998.
- [75] Alejandro Bassi, Ennio A. Vivaldi, and Adrián Ocampo-Garcés. The time course of the probability of transition into and out of REM sleep. *Sleep*, 32(5):655–669, 2009.

- [76] Adrián Ocampo-Garcés, Alejandro Bassi, Enzo Brunetti, Jorge Estrada, and Ennio A Vivaldi. REM sleep-dependent short-term and long-term hourglass processes in the ultradian organization and recovery of REM sleep in the rat. *Sleep*, 43(8):zsaa023, 2020.
- [77] Wei Wang, Robin K Yuan, Jude F Mitchell, Kirsi-Marja Zitting, Melissa A St. Hilaire, James K Wyatt, Frank AJL Scheer, Kenneth P Wright Jr, Emery N Brown, Joseph M Ronda, et al. Desynchronizing the sleep–wake cycle from circadian timing to assess their separate contributions to physiology and behaviour and to estimate intrinsic circadian period. *Nature protocols*, 18(2):579–603, 2023.
- [78] Tanya Dimitrov, Mingjian He, Robert Stickgold, and Michael J. Prerau. Sleep spindles comprise a subset of a broader class of electroencephalogram events. *Sleep*, 44(9):zsab099, 2021.
- [79] Clifford B Saper, Patrick M Fuller, Nigel P Pedersen, Jun Lu, and Thomas E Scammell. Sleep state switching. *Neuron*, 68(6):1023–1042, 2010.
- [80] Hiroto Ito and Akihiro Yamanaka. Deficiency of orexin signaling during sleep is involved in abnormal REM sleep architecture in narcolepsy. *Proceedings of the National Academy of Sciences*, 120(42):e2301951120, 2023.
- [81] Milo D Cornelissen, Simone Pilon, Luuk Steemers, Martin J Wanner, Steven Frolke, Ed Zuidinga, Steen Ingemann Jorgensen, Jarl Ivar van der Vlugt, and Jan H van Maarseveen. A covalent and modular synthesis of homo-and hetero [n] rotaxanes. *The Journal of Organic Chemistry*, 85(5):3146–3159, 2020.
- [82] Gert Jan Lammers, Giuseppe Plazzi, Emmanuel Mignot, Fabio Pizza, Yves Dauvilliers, Lucie Barateau, Paul Maruff, Thomas E Scammell, Robert D Latzman, Melissa Naylor, et al. Effects of orexin, an orexin receptor 2-selective agonist, on cognition in narcolepsy type 1: A secondary analysis of a randomized clinical trial. *JAMA neurology*, 83(2):145–152, 2026.
- [83] Tran V Hoc and Hsin-Chien Lee. Clinical and polysomnographic characteristics of Asian patients with comorbid insomnia and obstructive sleep apnea. *Scientific Reports*, 15(1):11529, 2025.
- [84] Cristian Popovici, Paschalis Steiropoulos, Stefan Mihaicuta, Silvia Dumitru, Sofia Dragila, Andras Bikov, and Alexandru Corlateanu. The hidden burden of COMISA: Clinical implications and treatment challenges. *Pulmonary Therapy*, pages 1–14, 2025.
- [85] Umaer Hanif, Flynn Crosbie, Anis Aloulou, Francesco Romano, Fabien Sauvet, Poul Jennum, Geoffroy Solelhac, and Damien Leger. REM sleep fragmentation and arousability distinguish chronic insomnia subtypes and reflect subjective-objective sleep discrepancy. *Sleep Medicine*, page 109165, 2026.
- [86] HW Agnew Jr, Wilse B Webb, and Robert L Williams. The first night effect: An EEG study of sleep. *Psychophysiology*, 2(3):263–266, 1966.

Methods

Cohorts and data

The shared data substrate comprised four clinical and epidemiological PSG cohorts (Table 2). Raw signals were stored in a harmonized 128 Hz HDF5 representation, with channels retaining cohort- and site-specific names. A shared channel-group registry mapped physiological modalities to the corresponding channel names at each site, allowing downstream pipelines to operate across cohorts without hard-coded site-specific channel labels. Each case study applied its own eligibility, feature-availability, and quality-control criteria; the resulting analytical samples are therefore reported separately in the corresponding cohort-flow tables.

SSC. The Stanford Sleep Clinic (SSC) [8, 10] is a clinical PSG cohort acquired at Stanford University Medical Center. Linked electronic health records provide demographics, ICD-coded diagnoses, medications, BMI, and longitudinal follow-up for incident disease analyses.

HSP. The Human Sleep Project (HSP v2.0) [24] is a clinical PSG cohort available through the BIDMC Brain Data Science Portal (<https://bdsp.io/content/hsp/2.0>). Linked electronic health records provide demographics, diagnoses, and longitudinal diagnostic follow-up used for external validation.

BioSerenity. The BioSerenity clinical network [25, 26, 27] comprises PSG recordings collected across a multi-site US clinical setting. Available information includes respiratory measures, sleep questionnaires, comorbidity history, and self-reported sleep symptoms. The PSG substrate is distinct from the larger questionnaire database used as the starting source for the COMISA cohort flow.

CDH. The Central Disorders of Hypersomnolence (CDH) [28, 29] registry is a Stanford-curated multi-site cohort assembled from seven institutions across France, China, Italy, Austria, Korea, and Denmark. It includes harmonized diagnoses of narcolepsy type 1, narcolepsy type 2 (NT2), idiopathic hypersomnia, secondary hypersomnia, and matched controls.

Software and reproducibility

AI SLEEP CO-SCIENTIST was implemented in Python 3.11 using the Claude Agent SDK (v0.1.76) and the Claude Code command-line runtime (@anthropic-ai/claude-code). All language-model calls used Claude Opus 4.6 (claude-opus-4-6). Agents executed inside a Singularity container built from a version-controlled definition file (python_node_3.11_22.def) containing Python 3.11, Node 22, the statistical analysis stack (lifelines, statsmodels, scikit-learn, numpy, scipy, and mne), and the version-pinned feature-extraction tools required by the relevant pipelines (Supplementary Table 1). Jobs were submitted to a shared HPC cluster using the same container environment for development, sample-scale validation, and full-cohort processing. Each run generated a versioned workspace containing the executed scripts, tool logs, sub-agent prompts, iteration summaries, structured critique verdicts, and resulting analysis artefacts. These workspaces preserve the analytical trajectory and provide the primary reproducibility artefacts for each case study. Because language-model generation is stochastic, independently rerunning an agent is not expected to reproduce the same trajectory exactly; instead, the reported numerical results are reproduced from the preserved analysis scripts and fixed data inputs.

System architecture

AI SLEEP CO-SCIENTIST comprises an Interface Layer and three specialist agents that can be invoked individually or iteratively as a research question develops. The Interface Layer provides the conversational surface through which the expert directs work, launches agents, reviews outputs, and supplies feedback. The Hypothesis Agent converts a research direction into a scored portfolio of data-grounded candidate questions. The Preprocessing Agent develops and validates feature-extraction pipelines from raw PSG signals on sample recordings. The Execution Agent carries a selected hypothesis through statistical analysis to a draft manuscript with code-grounded results. The three specialist agents are launched and reviewed through the Interface Layer, which maintains the persistent conversational state for each project; their analytical artefacts and intermediate outputs persist in versioned workspaces.

Three design commitments span these components. First, filesystem-based continuity: successive specialist-agent launches do not depend on persistent model context, but exchange state through versioned files on shared storage. This makes interrupted or iterative runs resumable and preserves an audit trail of prompts, tool calls, outputs, and feedback. Second, code-grounding: each quantitative result reported by the Execution Agent is linked to the executed analysis script, input data product, and output artefact from which it was derived. Third, expert-controlled checkpoints: in the expert-guided configuration evaluated here, the agents autonomously develop sample-scale preprocessing pipelines, execute analyses, and draft reports, while the expert approves full-cohort feature extraction, selects hypotheses for execution, and determines whether a manuscript is ready for finalization.

Shared infrastructure

All four components operate against a shared substrate that defines which data are available and how they can be accessed. Preprocessed feature sources are exposed through a uniform *describe*, *list*, and *load* interface implemented as a Model Context Protocol (MCP) server. Agents therefore inspect and access available data through schema-validated calls rather than relying on prompt-embedded file documentation that may become inconsistent with the underlying data. A canonical data-split function provides the authoritative top-level assignment of recordings to training, validation, and test sets, where applicable. Downstream analyses may construct folds within the training partition but cannot redefine validation or test membership. This keeps reported performance comparable across runs and guards against data leakage, which an agent free to define its own partition could otherwise introduce. The split definition is mounted separately from the selected cohort, and the container does not launch if this mount is absent. A shared skills library provides reusable procedural guidance and enforced figure-style standards, including colour-blind-safe palettes, typography, and axis-labelling conventions, thereby providing a common source of visual and methodological guidance across agents.

Each specialist agent runs inside a Singularity container with deliberately asymmetric filesystem permissions. System code, shared libraries, cohort features, raw signals, the hypothesis specification, and expert-feedback files are mounted read-only. The per-experiment workspace is writable and contains the analysis code, intermediate results, figures, reports, and logs generated during the run. These permissions are enforced by the container rather than by prompt instructions, preventing an agent from modifying source cohort data, authoritative split assignments, or its governing research specification. Session logs and machine-readable records of tool calls, including their inputs, outputs, and issuing agent, are written automatically to the experiment workspace, allowing actions by lead and sub-agents to be distinguished during post-hoc audit.

Interface Layer

The Interface Layer is a FastAPI web application through which the expert directs the system, reviews outputs, and provides feedback. It is the only component that maintains persistent conversational state for each project across sessions; specialist-agent state is instead preserved through the versioned workspace artefacts described above.

The interface serves three roles. First, it acts as a thinking partner during question formulation, helping the expert relate the scientific objective to the available cohorts and determine what evidence would constitute a convincing result. These discussions are grounded in the available data through the shared MCP interface. Second, it launches and monitors specialist agents through an MCP orchestration server that manages project lifecycle, hypothesis curation, Slurm job submission, status polling, and feedback authoring. Server-side validation restricts launches to supported cohorts and launch modes; in the expert-guided workflow evaluated here, launch parameters are displayed for confirmation before a cluster job is submitted. Third, it provides an audit and feedback channel. The interface reads specialist-agent artefacts (critique verdicts, cohort-flow records, iteration summaries, figures, and manuscript drafts) and converts expert comments into numbered, file-anchored actions while preserving the original feedback verbatim. It then supplies this feedback file to the relevant specialist agent, which treats it as the highest-priority input during the subsequent iteration.

The Interface Layer maintains one persistent Claude Agent SDK session per project. The session is reused across chat turns so that prior messages, file reads, and tool calls remain available without being reintroduced manually. When a project session is initialized, its system prompt is assembled from a static role specification, project metadata, and a project-memory file (`MEMORY.md`). The interface updates this file periodically as a concise record of project status and outstanding decisions rather than as a complete transcript. The SDK session store preserves conversational continuity across server restarts by replaying the stored session, while the project workspace preserves the corresponding analytical artefacts.

Hypothesis Agent

Given a research direction and one or more target cohorts, the Hypothesis Agent produces a scored portfolio of candidate scientific questions while preserving the lineage of each candidate across refinement steps. It runs parallel refinement trajectories within predefined diversity bins spanning scientific domains and eligible cohort contexts. At each step, several variants are generated from the current candidate and the highest-scoring variant is retained for further refinement. Four sub-agents operate under a lead orchestrator: a knowledge aggregator that assembles relevant literature and prior hypotheses; a generator that proposes candidate questions using this context and web-based literature retrieval; a critic that evaluates candidates against a multidimensional rubric; and a portfolio manager that records selected and rejected candidates, maintains the search graph, and assembles the final portfolio.

Two gates are applied before a candidate enters the search graph. The first assesses data feasibility. Variables named in a hypothesis are checked through the shared MCP interface against the cohort data catalogue, including preprocessed physiological features, demographics, clinical metadata, and available outcomes. If a required physiological feature is unavailable, the raw-signal channel registry is queried to determine whether it could be derived through a new preprocessing pipeline. Candidates requiring variables that are neither available nor derivable are rejected before scoring. The second gate audits the candidate’s scientific-question contract. Agent-generated candidates that assert anticipated effect sizes or sample sizes, or unnecessarily prescribe statistical methods and covariates, receive a capped score so that unsupported analytical detail is not mistaken for scientific specificity. This restriction applies to candidate generation; an expert may subsequently provide prespecified methods, covariates, or other requirements as part of the execution contract.

Each candidate is independently evaluated by the critic at two sampling temperatures. Disagreement between the resulting scores reduces that candidate’s priority during parent selection, limiting the influence of unstable evaluations. Search continues until the predefined round budget is exhausted or the trajectories meet stopping

criteria for score saturation or loss of portfolio diversity. The resulting high-scoring, diverse portfolio is presented through the Interface Layer. In the expert-guided studies reported here, candidates were selected for execution by the expert rather than promoted automatically.

Preprocessing Agent

When a hypothesis requires a physiological feature that is not already available in the shared data substrate, the Preprocessing Agent develops and validates the corresponding extraction pipeline. Given a free-text feature specification and one or more target cohorts, it implements a generalizable pipeline and evaluates it on a small sample of recordings before any full-cohort processing is initiated. Its initial deliverables comprise sample-level feature outputs, a rerunnable extraction script, methods documentation, and an expert-handoff report describing the changes required to expose the feature to downstream agents. Full-cohort feature extraction is initiated only after expert review and approval.

The Preprocessing Agent comprises a lead orchestrator and four sub-agents: a researcher responsible for literature and tool discovery and raw-signal inspection; an executor that implements and tests the pipeline; a reporter that prepares the expert-handoff documentation; and a critic that reviews correctness, physiological validity, and generalizability. The lead coordinates the workflow and routes tasks and artefacts among these specialist roles.

Five principles govern preprocessing. First, execution is restricted initially to the sample defined in the feature specification, typically tens of recordings. Pipelines are parameterized so that the same extraction code can subsequently be applied to the full cohort by changing the sample restriction after expert approval. This sample-first design allows methodological and computational problems to be identified before committing substantial cluster resources. Second, pipelines cannot hard-code channel names. Each script queries the shared channel-group registry at runtime and intersects the registered modality mappings with the channels present in each recording, allowing site-specific naming differences to be handled without changes to the extraction logic. Third, the researcher evaluates established community tools before proposing a new implementation. When an appropriate published tool is available, reuse provides a validated methodological starting point, although its behaviour is still evaluated on the target cohort. Fourth, pipelines parallelize processing across recordings to support full-cohort execution on the HPC cluster. Fifth, the critic reviews each iteration for hard-coded channels, cohort-specific assumptions, missing-signal handling, and other implementation choices that could prevent generalization; unresolved violations are classified as critical issues.

In the final iteration, the reporter produces an expert-handoff document containing, for each proposed feature, the tool and version used, input channels, operational definition, units, output schema, sample-level distributions, quality-control results, and known limitations. It also includes validation figures and a checklist of properties requiring expert confirmation before full-cohort extraction. After approval and full-cohort processing, the resulting feature source is registered in the shared data-access library and exposed to downstream agents through the same MCP interface. The complete library therefore combines investigator-provided features and pipelines with features added through the Preprocessing Agent, as summarized in Supplementary Table 1.

Execution Agent

The Execution Agent carries a selected hypothesis through staged analysis to a draft manuscript with code-grounded results. Each quantitative claim included in the manuscript must be supported by a stored result artefact and linked to the Python script that generated it. The scripts, inputs, outputs, figures, and tables are retained in the experiment workspace for audit and, subject to data-access requirements, independent re-execution. The agent comprises a lead orchestrator and four sub-agents: a researcher responsible for literature review and dataset inspection; an executor that writes and runs the stage-specific analysis scripts; a reporter that drafts the manuscript during the final iteration; and a critic that independently reviews the analytical record and returns a structured verdict. Each sub-agent has a role-specific tool allowlist. In particular, the critic can inspect the project artefacts but cannot execute shell commands or modify files, structurally separating review from the generation of code and results.

Hypothesis and execution specification as contract. A central principle of the Execution Agent is that the selected hypothesis and its accompanying expert-provided specification jointly define the minimum required analysis. Concrete requirements, including the cohort, sample construction, model, covariates, outputs, and literature to consider, must be followed unless they are infeasible or the expert explicitly approves a change. The agent may add sensitivity analyses, alternative models, or robustness checks, but

cannot silently omit or substitute a requested analysis. This prevents, for example, replacing a specified method with one judged to be “approximately equivalent” or analysing a smaller cohort because it is more convenient. Three mechanisms enforce this contract. First, required instructions are passed verbatim to each relevant sub-agent rather than only through an orchestrator’s paraphrase, reducing the risk that methodological constraints are lost during delegation. Second, when the agent identifies a methodological concern, it retains the specified analysis when it remains statistically and computationally defined, clearly reports the concern, and performs the proposed alternative as an additional analysis for expert adjudication. Third, the agent halts or defers a required step only when it cannot be executed with the available data, tools, or computing environment, for example because a required file, variable, or software dependency is unavailable. In such cases, it records the diagnostic evidence and reports the unresolved requirement rather than substituting a different procedure.

Code-grounding. The Execution Agent requires each analysis workspace to contain a standard set of structured provenance artefacts. A fixed directory structure separates scripts, results, figures, and logs. Each analytical stage has a corresponding script with a machine-readable header documenting its purpose, inputs, outputs, parameters, and runtime. Before figure generation, a figures-and-tables plan maps each intended output to the script and data products that will generate it; subsequent changes to the plan are retained in the workspace history. A features ledger records each covariate and engineered variable together with its operational definition, units, plausible range where defined, and supporting literature where relevant. A methods ledger records analytical decisions, including their rationale, parameters, assumptions assessed, and sample sizes, which the reporter uses when drafting the Methods section. These ledgers reduce the risk that undocumented analytical choices are later described as prespecified or validated. A cohort-flow artefact records the starting population, sequential exclusions and their rationales, and the final analytical sample. Its arithmetic is checked mechanically, and it is rendered as the first cohort table in the internal case-study report. Across iterations, the critic compares successive cohort-flow versions and flags unexplained changes in sample construction. During the final iteration, the reporter creates a unified reproduction script that invokes the staged analysis code to regenerate the reported numerical results, figures, and tables from the fixed analytical inputs. This script provides a single entry point for reproducing the manuscript outputs; full raw-signal feature extraction remains a separate preprocessing step when rerunning it would be computationally prohibitive.

Automated quality control and hallucination mitigation. Beyond code-grounding, the Execution Agent applies automated checks for several common failure modes of language-model-generated scientific reports. Each citation is entered into a structured registry containing its title and source URL, allowing the reference itself to be traced to the retrieval event from which it originated. References whose bibliographic identity cannot be verified are flagged and excluded from the final agent-generated manuscript rather than completed from model memory. This check addresses fabricated or incorrectly specified references but does not, by itself, establish that a cited source supports the associated interpretation. Figure-style compliance is assessed in both the plotting code and the rendered output. Plotting scripts are checked for use of the shared figure-style package, while rendered PDFs are inspected for violations of the required palette, typography, labelling, and layout conventions. After manuscript drafting, standalone-readability checks flag workspace-specific paths, internal variable or feature names, and undefined acronyms that would prevent the report from being interpreted independently of its analysis workspace. These automated checks reduce identifiable errors but do not establish scientific validity, which remains subject to expert review.

Iteration model and expert feedback. At the end of each iteration, the critic returns a structured verdict comprising an overall quality label and counts of critical, important, and optional issues. When expert feedback has been supplied, any feedback item that remains unresolved is classified as critical and limits the overall verdict to “Needs Improvement”. The orchestrator terminates the refinement loop when the critic reports no critical issues, when the configured iteration budget is exhausted, or when two consecutive iterations satisfy the predefined stagnation criterion. When the Interface Layer relaunched the Execution Agent with expert feedback, the lead creates an itemized response file that preserves each comment verbatim and records the corresponding required action, implementation status, and affected code, result, figure, or manuscript location. Relevant executor prompts include both the original comment and its translated action, reducing the risk that constraints are lost through paraphrase. Before invoking the critic, the lead marks each item as completed or deferred and provides supporting evidence or a rationale. The critic then independently inspects the cited artefacts and verifies whether each requested change has been implemented. Deferred items remain unresolved, and therefore critical, unless the expert explicitly accepts the deferral. Feedback iterations are non-destructive. Before a new iteration begins, the previous report, execution artefacts, and ledgers are copied to numbered version snapshots, allowing successive analyses and drafts to be compared. At the end of

each iteration, the lead writes a structured summary of the completed work, unresolved issues, and current artefact locations. These summaries provide the primary filesystem-based memory for refinement across specialist-agent launches without requiring the complete prior conversation to be reloaded.

Data availability

HSP v2.0 is distributed through the BIDMC Brain Data Science Portal (<https://bdsp.io/content/hsp/2.0>). The CDH cohort is available at <https://stanfordmedicine.app.box.com/s/r9e92ygg0erf7hn5re6j51aaggf50jly>, and SSC is available through the Brain Data Science Portal at <https://bdsp.io/content/08vg8vqv2wdtwonc1ddy/1.0/>. Access to linked electronic health records, longitudinal outcomes, and other restricted clinical variables is governed by cohort-specific approvals and may differ from the data included in the public PSG releases.

The BioSerenity data are proprietary and cannot be redistributed publicly. Requests for access are subject to approval by BioSerenity and the applicable institutional, privacy, and data-use requirements.

A Supplementary Material

A.1 Preprocessed feature library (extended)

The shared feature library exposes preprocessed representations of raw PSG signals and associated metadata through a uniform Model Context Protocol (MCP) interface (see Methods). Feature availability across the four cohorts is summarized in Supplementary Table 1. Features are grouped by pipeline origin. *Human-led* denotes features or pipelines selected, supplied, or previously established by investigators, including curated clinical metadata and outputs from existing preprocessing systems. *Agent-led* denotes pipelines curated, adapted, or implemented by the Preprocessing Agent under expert review; it does not imply that the underlying physiological method or software was invented by the agent. Feature definitions, output formats, and relevant cohort-specific considerations are described below.

A.1.1 Human-led features

Sleep stages (U-Sleep). U-Sleep [30] is a fully convolutional neural network that maps EEG and EOG signals to sleep-stage probabilities. Outputs are stored as *hypnodensities*: a $T \times 5$ probability matrix over Wake, N1, N2, N3, and REM at 1-second resolution. Hard-stage labels are derived by taking the highest-probability stage when required. Outputs are stored separately for each recording and available channel configuration, together with processing logs.

Cortical arousals (MAD). The Multimodal Arousal Detector (MAD) [31] identifies cortical arousals using available EEG, EOG, EMG, and ECG signals. Event-level outputs contain the onset, offset, duration, and predicted probability of each detected arousal.

Respiratory events (ABED). The Apnoea-Based Event Detector (ABED) [32] identifies obstructive and central apnoea and hypopnoea events from available respiratory-effort, airflow, and oximetry signals as well as MAD-based arousal detections. Event-level outputs contain onset and offset times, event type, and either the predicted probability or corresponding binary classification.

EEG spectrograms. Time–frequency spectrograms were precomputed from available EEG channels to avoid repeating spectral decomposition during downstream analyses. Spectrograms were generated using the LSP-optimal multitaper spectrogram with 30-second non-overlapping windows. Spectral power was converted to decibel units and only frequencies between 0.5–25 Hz were retained. Spectrograms are stored separately by recording and channel.

Transient oscillations (DynamO). The Dynamic Oscillation Toolbox (DynamO) [33] detects transient EEG oscillations, defined as brief, spectrally resolved bursts distributed across the time–frequency plane. For each channel, sleep stage, and frequency band, the stored outputs support six derived feature families: oscillation rate, prominence, duration, slow-oscillation phase-coupling strength, preferred slow-oscillation phase angle, and local slow-oscillation power.

Brain–heart coupling. Coupling between EEG and cardiac activity was computed for each available EEG channel using an established brain–heart interaction formulation [34]. Outputs are stored as per-recording, per-channel arrays for downstream feature summarization.

Demographics. Available subject-level metadata include age, sex, BMI, race, and ethnicity. Availability and coding differ across cohorts, and race/ethnicity follows each cohort’s native conventions (Table 2).

Diagnoses and curated outcomes. SSC and HSP contain linked ICD-9/ICD-10 codes and longitudinal time-to-diagnosis information. CDH contains harmonized primary and secondary diagnostic labels, including NT1, NT2, idiopathic hypersomnia, secondary hypersomnia, and controls. BioSerenity does not contain linked diagnostic codes in the shared substrate; the COMISA groups were instead derived from self-reported insomnia-related questionnaire and objective respiratory measures, as described in the corresponding case-study methods.

Medications. SSC contains longitudinal medication records extracted from linked electronic health records. Other cohorts either do not provide medication information or contain only limited, non-harmonized indicators; these were not exposed as a common medication feature source.

Sleep questionnaires. BioSerenity contains subject-reported sleep symptoms and behaviours, including the Epworth Sleepiness Scale and its component items, self-reported medical and problem history, and sleep-related symptoms and behaviours.

A.1.2 Agent-led features

R-peaks. R-peak locations were extracted from available ECG channels and represented as event times within each recording. These detections provide the input for downstream heart-rate-variability, brain–heart-coupling, and time-delay-stability analyses.

Spindles (YASA). Sleep spindles were detected during NREM sleep using YASA [35]. Event-level outputs include spindle onset, duration, peak frequency, amplitude, and the EEG channel from which the event was detected.

Slow waves (YASA). Slow waves were detected during NREM sleep using YASA [35]. Event-level outputs include onset, negative- and positive-peak timing, amplitude, duration, and the EEG channel from which the event was detected.

Eye movements (YASA). Eye-movement events were detected from available EOG channels using YASA [35]. Event-level outputs include onset, offset, duration, detection channel, and event classification where available.

EEG artefact masks (YASA). YASA [35] was used to identify EEG segments likely to contain artefact. The resulting time-aligned Boolean masks allow downstream pipelines to exclude contaminated periods from feature extraction and statistical summaries.

Continuous oximetry. Available pulse-oximetry signals were extracted into a standardized, time-aligned representation retaining the SpO₂ value and sampling information for each recording.

Desaturation events. Discrete oxygen-desaturation events were derived from the continuous SpO₂ signal. Event-level outputs include onset, offset, nadir, and the magnitude of the decrease from the preceding baseline.

Heart-rate variability. Heart-rate-variability features were calculated from the R-peak time series and stratified by sleep stage. The stored summaries include time-domain measures such as SDNN, RMSSD, and pNN50 and frequency-domain measures such as low-frequency power, high-frequency power, and their ratio.

Brain–EMG coupling. The brain–heart coupling framework was adapted to quantify stage-stratified coupling between EEG and available EMG signals. Outputs summarize cortical–muscular coupling for each available EEG–EMG channel pair.

Time-delay-stability networks. Time-delay stability (TDS) [36] quantifies coupling among available cortical, cardiac, and peripheral physiological signals. For each signal pair, sliding-window cross-correlation was used to estimate the lag at which coupling was maximal. TDS link strength was defined as the proportion of windows in which this lag remained stable. A signal pair was classified as a significant link when more than 7% of its windows met the stability criterion. Features were summarized by sleep stage and physiological subnetwork, including brain–brain, brain–heart, brain–body, and body–body coupling.

A.2 Additional Case Study Details

A.2.1 Transient-oscillation features

Transient oscillations (TOs) were detected with the Dynamic Oscillation Toolbox (DynamO) [33], which identifies them as *time-frequency peaks* (TF-peaks): salient, topographically distinct high-power regions in the multitaper spectrogram of the EEG (1 s windows, 0.05 s steps, time-half-bandwidth 2, 3 tapers). A watershed segmentation is applied to the negative of this topography; the resulting over-segmented regions are merged under a completeness/distinctness weight rule, each surviving region is trimmed to the boundary enclosing 80% of its volume, and peaks falling below the estimator’s own time (0.5 s) and frequency (2 Hz) resolution are discarded as noise. We used the toolbox’s recommended **fast** preset: 30 s segments, 2×2 topographic downsampling, and a merge threshold of 11, applied iteratively until no pair of regions exceeded it. Because peaks below 4 Hz cannot be disambiguated from the slow oscillation itself, analysis was restricted to 4–25 Hz and partitioned into five bands (θ 4–8, low- α 8–11, slow- σ 9–12, fast- σ 12–16, β 16–25 Hz).

Six features summarize the TOs within each channel $\times$ stage $\times$ band:

- **Rate:** number of TOs per unit time spent in that stage.
- **Log-prominence:** log of the TF-peak height above its local spectral baseline ($\mu\text{V}^2/\text{Hz}$), averaged over TOs; an amplitude measure.
- **Duration:** temporal width of the TF-peak (s), averaged over TOs.
- **SO-power:** slow-oscillation power (0.3–1.5 Hz, integrated and expressed in dB from a separate 30 s multitaper spectrogram) *at the time of each TO*, averaged over TOs. This indexes the slow-wave background in which the TOs occur, not stage-wide SO power.
- **SO-phase angle:** circular mean of the slow-oscillation phase at TO times. Phase is obtained by band-passing the EEG at 0.3–1.5 Hz with a zero-phase filter and taking the Hilbert angle, wrapped so that 0 rad is the SO peak and $\pm\pi$ the SO trough (surface-negative convention; Figure 6a).
- **SO-phase coupling strength (MRL):** mean resultant length of those same phases (0 = uniform across the SO cycle, 1 = perfectly phase-locked); it measures the *concentration* of coupling, independently of the preferred angle.

Rate, log-prominence, duration and SO-power were residualized against age, sex, BMI and site (rate additionally against stage proportions); the two circular measures are not amenable to ordinary least-squares residualization and are reported as raw circular statistics. The scan therefore comprises 6 features $\times$ 6 EEG channels (F3, F4, C3, C4, O1, O2) $\times$ 5 sleep stages (Wake, N1, N2, N3, REM) $\times$ 5 frequency bands = **900 features**, each contributing one NT1-versus-control Cohen’s d and one Benjamini–Hochberg FDR-corrected q value (Supplementary Figure 14).

Table 1: **Preprocessed feature library across cohorts.** Feature families exposed through the shared data substrate, grouped by pipeline origin. Human-led denotes features or pipelines selected or provided by investigators; agent-led denotes pipelines curated, adapted, or implemented by the Preprocessing Agent under expert review and does not imply that the underlying method was invented by the agent. ✓ indicates availability for at least a subset of eligible recordings; – indicates that the feature is unavailable in the current substrate. Recording-level availability may vary because of missing signals and feature-specific quality control. SSC = Stanford Sleep Clinic; HSP = Human Sleep Project [24]; CDH = Central Disorders of Hypersomnolence.

Origin	Feature	SSC	HSP	BioSerenity	CDH
Human-led	Sleep stages (U-Sleep [30])	✓	✓	✓	✓
	Cortical arousals (MAD [31])	✓	✓	✓	✓
	Respiratory events (ABED [32])	✓	✓	✓	✓
	EEG spectrograms	✓	✓	✓	✓
	Transient oscillations (DynamO [33])	–	–	–	✓
	Brain–heart coupling [34]	✓	✓	✓	✓
	Demographics	✓	✓	✓	✓
	Diagnoses (ICD or curated outcomes)	✓	✓	–	✓
	Medications	✓	–	–	–
	Sleep questionnaires	–	–	✓	–
Agent-led	R-peaks	✓	✓	✓	✓
	Spindles (YASA [35])	✓	✓	✓	✓
	Slow waves (YASA [35])	✓	✓	✓	✓
	Eye movements (YASA [35])	✓	✓	✓	✓
	EEG artefact masks (YASA [35])	✓	✓	✓	✓
	Continuous oximetry features	✓	✓	✓	✓
	Desaturation events	✓	✓	✓	✓
	Heart-rate variability	✓	✓	✓	–
	Brain–EMG coupling	✓	✓	–	–
	Time-delay-stability networks (TDS [36])	✓	✓	–	–

Table 2: **Characteristics of the four PSG cohorts represented in the shared data substrate.** N denotes the number of PSG recordings in each cohort-level snapshot before case-study-specific eligibility, feature-availability, and quality-control criteria were applied. Continuous variables are reported as mean $\pm$ SD among recordings with the corresponding variable available, and sex and race/ethnicity as percentages. Race/ethnicity follows cohort-native coding and may not sum to 100% because of missing or unreported values. The analytical samples used in individual case studies are reported separately in the corresponding cohort-flow tables.

Variable	SSC	HSP	BioSerenity	CDH
N (PSG recordings)	35,018	19,590	67,830	1,832
Age, mean $\pm$ SD (years)	45.5 $\pm$ 19.8	51.9 $\pm$ 16.6	51.2 $\pm$ 18.1	31.1 $\pm$ 15.9
Age range (years)	0–99	1–100	1–103	3–90
Female sex (%)	40.6%	43.3%	48.5%	42.9%
BMI, mean $\pm$ SD (kg/m ²)	28.2 $\pm$ 7.2	–	33.1 $\pm$ 9.1	25.3 $\pm$ 5.6
<i>Race/ethnicity (%; cohort-native coding)</i>				
White	54.0%	77.6%	69.3%	–
Asian	12.9%	3.4%	0.6%	–
African American	2.9%	6.2%	23.1%	–
Hispanic	9.2%	–	2.8%	–
Other	21.0%	10.2%	0.9%	–

Step 1. The script declares its inputs, outputs, and runtime in a machine-readable header.

```

18 Inputs:
19 - /workspace/execution/results/features_ssc.parquet (28,817 × 421)
20 - /workspace/execution/results/features_s0001.parquet (16,973 × 421)
21 - /workspace/execution/results/feature_dictionary.csv (397 features)
22
23 Outputs:
24 - /workspace/execution/models/clock_<c>.pkl (8 clocks)
25 - /workspace/execution/models/agg_global.pkl
26 - /workspace/execution/models/agg_combined_meta.pkl
27 - /workspace/execution/results/age_clock_performance.csv
28 - /workspace/execution/results/age_clock_stratified_performance.csv
29 - /workspace/execution/results/predictions_ssc.parquet
30 - /workspace/execution/results/predictions_s0001.parquet
31 - /workspace/execution/results/residuals_ssc.parquet
32 - /workspace/execution/results/residuals_s0001.parquet
33 - /workspace/execution/results/bias_correction_params.csv
34 - /workspace/execution/results/residual_structure.csv
35 - /workspace/execution/results/hrv_clip_caps.json
36 - /workspace/execution/results/ahi_sensitivity_clocks.csv
37 - /workspace/execution/results/stacker_bootstrap_stability.csv
38 - /workspace/execution/results/common_subject_performance.csv
39 - /workspace/execution/figures/fig03_age_clock_grid.pdf
40 - /workspace/execution/figures/fig04_residual_vs_age_grid.pdf
41
42 Runtime: ~20-30 min.
    
```

Step 2. All performance metrics come from a single helper.

```

186 def perf_metrics(y_true: np.ndarray, y_pred: np.ndarray) -> Dict[str, float]:
187     mask = ~(np.isnan(y_true) | np.isnan(y_pred))
188     y_true, y_pred = y_true[mask], y_pred[mask]
189     if len(y_true) < 5:
190         return dict(n=int(len(y_true)), mae=float("nan"), rmse=float("nan"),
191                    r2=float("nan"), slope=float("nan"), intercept=float("nan"))
192     mae = mean_absolute_error(y_true, y_pred)
193     rmse = float(np.sqrt(mean_squared_error(y_true, y_pred)))
194     r2 = r2_score(y_true, y_pred)
195     slope, intercept = calibration_slope_intercept(y_true, y_pred)
196     return dict(n=int(len(y_true)), mae=float(mae), rmse=rmse, r2=float(r2),
197               slope=slope, intercept=intercept)
    
```

Step 3. Each row is labelled with the model that produced it.

```

514 perf_rows = []
515 for cohort, y, pred in [
516     ("ssc_train", y_train, p_train),
517     ("ssc_val", y_val, p_val),
518     ("ssc_test", y_test, p_test),
519 ]:
520     m = perf_metrics(y, pred)
521     m.update(clock="combined_meta", cohort=cohort, estimator="ridge_stacker")
522     perf_rows.append(m)
    
```

Step 4. The assembled table is written to the result file named in the header.

```

1402 perf_df.to_csv(RESULTS_DIR / "age_clock_performance.csv", index=False)
    
```

Step 5. The reported value is read from that file, at full precision.

```

clock,cohort,estimator,n,mae,rmse,r2,...
combined_meta,ssc_test,ridge_stacker,4096,7.057098180235453,9.044013335769586,0.6860743374930721,...
    
```

Figure 7: Code-grounding of a reported quantity: multidomain sleep ageing. The main text reports an MAE of 7.06 years and $R^2 = 0.686$ for the multidomain late-fusion model on the held-out SSC test set. The steps trace that quantity through `stage2_iter7_train_clocks.py`: the declared inputs and outputs, the single helper that computes all performance metrics, the call that labels each row with the model producing it, the line that writes the result file, and the row from which the reported value is taken at full precision. Line numbers are those of the original script in the preserved workspace; excerpts are contiguous, indentation is reduced, and the result file is truncated at the quoted columns.

Step 1. The header declares inputs, outputs, and the parameters governing the analysis.

```

4 Purpose: Cox proportional-hazards (incident cases) and logistic regression (all cases)
5   to quantify the risk associated with lower TDS coupling for each comorbidity.
6 Inputs:
7   - /workspace/execution/results/tds_features_ssc.parquet: per-subject TDS features (33,202 x 397)
8   - /workspace/execution/results/cohort_ssc.csv: cohort with group assignments (35,038 rows)
9   - /workspace/execution/results/cohort_summary.json: summary statistics
10  - Expert-curated outcome data via get_SSC_stanford_outcome_df()
11  - Phecode outcome data via get_SSC_stanford_phecode_outcome_df()
12 Outputs:
13  - /workspace/execution/results/risk_analysis_hr.csv: all HR results
14  - /workspace/execution/results/risk_analysis_or.csv: all OR results
15  - /workspace/execution/results/risk_analysis.csv: combined summary
16  - /workspace/execution/figures/fig06_hazard_ratios.pdf: forest plot of HRs
17  - /workspace/execution/figures/fig07_odds_ratios.pdf: forest plot of ORs
18 Parameters:
19  - penalizer: 0.01 -- regularization for CoxPHFitter to aid convergence
20  - FDR q-threshold: 0.05 -- Benjamini-Hochberg correction within each comorbidity
21  - BMI imputation: median -- for missing BMI values (13% of SSC)
22  - Direction: predictors are NEGATED (per-SD decrease) so HR>1 = lower coupling -> higher risk
    
```

Step 2. A model is fitted for every predictor and sleep stage in turn.

```

213 for pred_base, pred_display in PREDICTORS.items():
214     for stage in STAGES:
215         col = f"{pred_base}_{stage}"
216         if col not in cox_data.columns:
217             continue
218
219         # Build Cox dataframe
220         cox_df = cox_data[[col, 'age', 'sex', 'bmi', 'time_to_event', 'label']].dropna()
    
```

Step 3. The predictor is standardized and negated.

```

231 # NEGATE: we want HR per SD DECREASE in coupling
232 # If we negate the z-scored predictor, HR > 1 means lower coupling -> higher risk
233 cox_df['tds_predictor'] = -((cox_df[col] - pred_mean) / pred_std)
234
235 cox_df_fit = cox_df[['tds_predictor', 'age', 'sex', 'bmi', 'time_to_event', 'label']].copy()
236 cox_df_fit.columns = ['tds_predictor', 'age', 'sex', 'bmi', 'T', 'E']
237
238 try:
239     cph = CoxPHFitter(penalizer=PENALIZER)
240     cph.fit(cox_df_fit, duration_col='T', event_col='E')
241
242     hr = np.exp(cph.params_['tds_predictor'])
243     ci = np.exp(cph.confidence_intervals_.loc['tds_predictor'].values)
    
```

Step 4. Each fitted model contributes one row, and the table is written to the result file.

```

258 hr_results.append({
259     'comorbidity': grp_name,
260     'display_name': display,
261     'predictor': pred_base,
262     'predictor_display': pred_display,
263     'stage': stage,
264     'feature': col,
265     'network': NETWORK_DISPLAY[pred_base],
266     'HR': hr,
267     'HR_lower': hr_lower,
268     'HR_upper': hr_upper,
269     'p_value': p_val,
270     'n_total': len(cox_df_fit),
271     'n_events': int(cox_df_fit['E'].sum()),
272     'ph_assumption_ok': ph_ok,
273     'pred_mean': pred_mean,
274     'pred_std': pred_std,
275 })
399 hr_df.to_csv(OUT_HR, index=False)
    
```

Step 5. The reported hazard ratio is read from that file, at full precision.

```

comorbidity,predictor,stage,HR,HR_lower,HR_upper
Parkinsons,mean_ls_brain_brain,N2,1.4778671263332086,1.3066306803136443,1.6715444356259153
    
```

Figure 8: **Code-grounding of a reported quantity: network-physiology coupling.** The main text reports a hazard ratio of 1.48 (95% confidence interval 1.31 to 1.67) for subsequent Parkinson’s disease per 1 SD decrease in N2 brain–brain coupling. The steps trace that quantity through `stage4_risk_analysis.py`. Line numbers are those of the original script; excerpts are contiguous, indentation is reduced, and the result file is shown for the quoted columns.

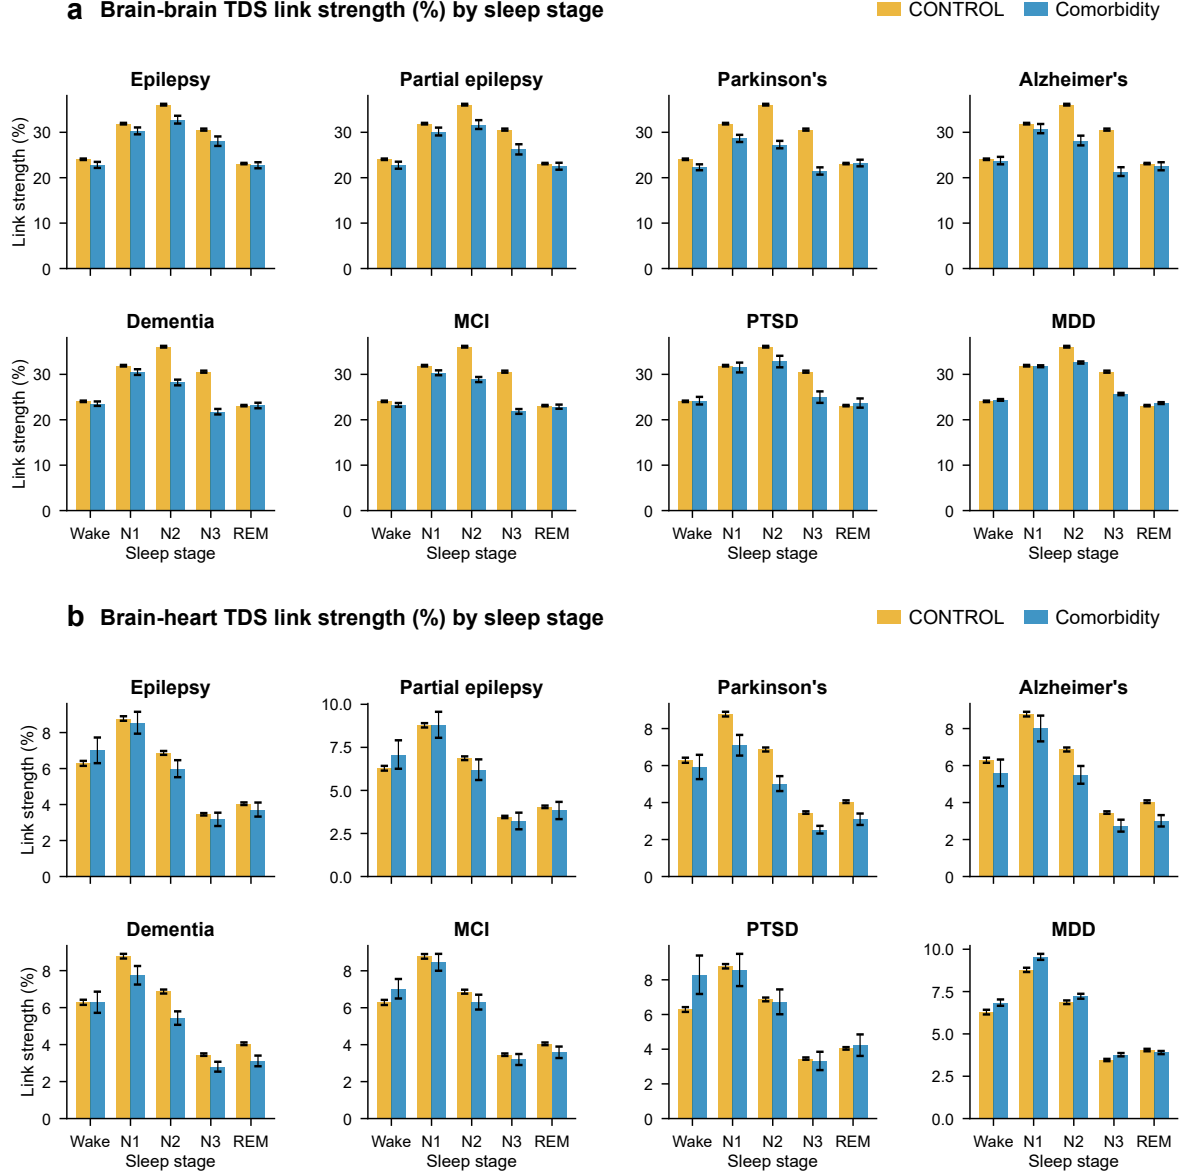


Figure 9: **Sleep-stage profiles of TDS link strength by disorder group in SSC.** (a) Brain–brain and (b) brain–heart TDS link strength (%) across sleep stages (Wake, N1, N2, N3, and REM), shown for SSC controls and each of the eight disorder groups: epilepsy, partial epilepsy, Parkinson’s disease, Alzheimer’s disease, dementia, MCI, PTSD, and MDD. Bars show unadjusted group means, and error bars show 95% confidence intervals. Each panel displays one disorder group alongside the SSC Control group; covariate-adjusted comparisons are presented in Supplementary Figure 10.




Figure 10: **Covariate-adjusted TDS coupling differences between disorder groups and controls in SSC.** Volcano plots show Hedges' g , calculated from feature residuals after adjustment for age, sex, and BMI, against FDR-adjusted statistical significance ($-\log_{10} q$) for each of the eight disorder groups. Positive g indicates lower coupling in the disorder group than in controls, whereas negative g indicates higher coupling. Each point represents a TDS coupling feature and is coloured by subnetwork (brain–brain, brain–heart, brain–body, body–body, or all-network summary). The dashed line marks the Benjamini–Hochberg FDR threshold ($q = 0.05$). Each panel reports $n_{\max}$, the largest per-feature analytic sample for that group, and the number of FDR-significant features; the analytic sample varies across features according to feature and covariate availability. The two labelled features in each panel have the smallest FDR-adjusted q -values and are not necessarily those with the largest absolute effect sizes. LS, link strength; %sig, percentage of significant links; BB, brain–brain; BH, brain–heart; B-EMG, brain–EMG; B-Eye, brain–eye; F, frontal; F-C, frontal–central.




Table 3: **Cohort flow and disorder-group counts for SSC and HSP.** (a) Sequential exclusions used to derive the post-TDS-QC cohorts. All counts represent unique subjects. (b) PheCode-defined group counts before and after the common TDS quality filter. Cohort n denotes subjects meeting the group definition after demographic eligibility but before TDS quality control; post-QC n denotes subjects who additionally passed the TDS quality filter. Feature-specific analytic samples may be smaller because of missing physiological signals. [†]The Control group required the absence of documented neurological conditions (epilepsy, Parkinson’s disease, dementia, Alzheimer’s disease, and MCI); neuropsychiatric conditions (MDD and PTSD); cardiovascular conditions (hypertension, atrial fibrillation, heart failure, acute myocardial infarction, and ischaemic heart disease); sleep disorders (insomnia and REM sleep behaviour disorder); and stroke or TIA. [‡]Generalized convulsive epilepsy (PheCode 345.11) and convulsions (PheCode 345.3) were excluded from the epilepsy and partial-epilepsy groups. The generalized-convulsive-epilepsy exclusion affected 408 SSC and 232 HSP subjects. Disorder groups were evaluated independently and could overlap, except that subjects with both MDD and PTSD were excluded from the individual MDD and PTSD comparisons and tracked separately. Dashes indicate groups that were not analysed separately after TDS quality control. HSP did not include BMI; HSP models were therefore adjusted for age and sex only.

a. Sequential cohort flow					
Step	Criterion	SSC (discovery)		HSP (replication)	
		Excluded	n after	Excluded	n after
0	Starting subjects	–	39,920	–	17,977
1	Precomputed TDSpy output available	4,868	35,052	0	17,977
2	Age and sex available	14	35,038	0	17,977
3	TDS feature-loading quality filter	1,836	33,202	167	17,810
Final post-QC cohort			33,202		17,810
b. PheCode-defined group counts					
Group		SSC (discovery)		HSP (replication)	
		Cohort n	Post-QC n	Cohort n	Post-QC n
Control [†]		11,809	11,240	3,632	3,604
Epilepsy [‡]		688	678	420	412
Partial epilepsy [‡]		500	494	334	328
Parkinson’s disease		629	612	303	302
Alzheimer’s disease		361	352	432	429
Dementia		1,006	986	966	958
MCI		1,242	1,207	1,030	1,019
PTSD		280	268	244	242
MDD		7,417	7,018	4,965	4,923
MDD+PTSD, excluded from both comparisons		921	–	1,055	–

Table 4: **Stage-1 quality cohort flow for SSC (development) and HSP v2.0 (external validation).** Sequential exclusion counts through the pre-specified QC pipeline: adult recordings (≥ 18 y) with known demographics, at least one hypnodensity channel, recording duration > 4 h, and total sleep time (TST) > 90 min. TST is computed from the hypnodensity/hypnogram. Recordings that pass stage-1 QC are further excluded at the feature-extraction stage where modality-specific features are unavailable, so the analytic samples reported elsewhere are smaller than the split sizes tabulated here. The SSC test split contains 4,120 recordings, of which 4,096 carried the features required for the reported multidomain late-fusion evaluation (Figure 3a); the analytic n for each individual domain clock is annotated in Supplementary Figure 11a.

QC step	SSC		HSP v2.0	
	n excluded	n remaining	n excluded	n remaining
Starting cohort (canonical split)	–	35,018	–	19,590
Age known	3	35,015	1,613	17,977
Age ≥ 18	3,799	31,216	427	17,550
Sex coded	10	31,206	0	17,550
Hypnodensity present (≥ 1 channel)	1,787	29,419	209	17,341
Recording duration > 4 h	322	29,097	258	17,083
Total sleep time > 90 min	280	28,817	110	16,973
Stage-1 quality cohort (final)		28,817		16,973
<i>Canonical split assignments (subjects entering downstream modelling):</i>				
Train		24,098		8,932
Validation		599		1,268
Test		4,120		6,773

Table 5: **Robustness of the sleep-ageing findings across the original analysis and three independent AI SLEEP CO-SCIENTIST reruns.** To test whether the principal conclusions depended on a particular agent trajectory, the complete study was rerun three times from the same high-level prompt, with each run independently constructing its own feature space and analysis workflow. For the original run and each rerun, the table reports the quantities most relevant to the study’s conclusions, together with an overall assessment. Values separated by “/” denote SSC and HSP, respectively. Feature counts are descriptive rather than directly comparable across runs, because each run independently engineered its feature space. Abbreviations: MAE, mean absolute error; HR, hazard ratio per 1 SD residual; R^2 , coefficient of determination; C-index, Harrell’s concordance index; PheWAS, phenome-wide association study. The cardiometabolic composite comprised acute myocardial infarction, heart failure, and stroke/TIA. PheWAS replication was evaluated among SSC FDR-significant incident PheCodes tested in HSP, reported as Spearman ρ and direction concordance. In the assessment column, “replicated” indicates the conclusion held across all runs, “partially replicated” the same qualitative pattern with quantitative variation, and “negative result replicated” that residual-only disease tuning failed to improve discrimination in every run. The Rerun 3 SSC composite HR was approximated from its forest plot.

Measure	Original	Rerun 1	Rerun 2	Rerun 3	Assessment
QC-passing recordings, SSC / HSP	28,817 / 16,973	28,543 / 16,928	28,571 / 16,928	28,341 / 16,906	Replicated
Engineered features (count)	362	622	1,105	365	Variable
Stacked-model age MAE, SSC / HSP (years)	7.06 / 8.83	8.00 / 13.20	6.79 / 8.51	8.41 / 11.46	Partially replicated
Stacked-model age R^2 , SSC / HSP	0.686 / 0.494	0.55 / -0.12	0.71 / 0.53	0.56 / 0.10	Partially replicated
Cardiometabolic composite HR (per 1 SD), SSC / HSP	1.23 / 1.26	1.15 / 1.21	1.18 / 1.15	≈ 1.22 / 1.15	Replicated
Incident PheWAS replication, ρ / concordance	0.33 / 92.0%	0.60 / 95.0%	0.54 / 88.7%	0.55 / 88.4%	Replicated
Residual-only disease-tuned C-index, SSC / HSP	0.517 / 0.495	0.555 / 0.452	0.473 / 0.447	0.542 / 0.547	Negative result replicated

Table 6: **Cohort flow for the COMISA case study.** The BioSerenity source contained 540,941 clinical and sleep-questionnaire records; the associated PSG cohort comprised 67,830 recordings. All tabulated counts refer to PSG recordings. Two analytic samples were derived independently from this cohort using the same four-group definitions and excluding recordings with AHI 5–14. The microstructure screen (n_1) required hypnodensity, arousal, spindle, and slow-wave features, whereas the arousal-focused analysis (n_2) required arousal, respiratory-event, and hypnodensity data. The first exclusion count in each branch is calculated relative to the 67,830-recording PSG cohort.

Criterion	Excluded	Remaining
BioSerenity PSG cohort	–	67,830
Microstructure screen (n_1)		
Four-group eligibility and all four microstructure domains	34,510	33,320
Non-empty feature files and ≥ 10 min sleep	1,482	31,838
Plausible covariates	27	31,811
Arousal-focused analysis (n_2)		
Four-group eligibility and required arousal, respiratory-event, and hypnodensity data	31,282	36,548
Recording duration ≥ 3 h	155	36,393
At least 10 arousal events	356	36,037
Total sleep time ≥ 0.5 h	73	35,964




Figure 12: Supplementary arousal-dynamics analyses for the COMISA replication cohorts. (a) Four-class classification performance (macro-averaged AUC, test set) across feature sets ranging from demographics and macro-PSG baselines to single microstructure domains and combined feature sets, for random-forest and logistic-regression classifiers. (b) Position of the COMISA phenotype along the insomnia-OSA spectrum: (left) Mahalanobis distance of each recording to the OSA and insomnia centroids, coloured by hybrid score, and (right) distribution of the hybrid score for OSA, insomnia, and COMISA, showing the COMISA distribution shifted towards the OSA end. (c) Domain-level discrimination, shown as macro-averaged AUC for each microstructure domain combined with age and sex, for random-forest and logistic-regression classifiers. (d) Principal-component analysis of the arousal-dynamics feature space. The diagnostic groups overlap extensively, with PC1 (36.4% of variance) separating OSA from non-OSA recordings. Inset, the variance explained by the first ten principal components.




Figure 13: **Replication of homeostatic REM-pressure findings in CDH and SSC cohorts.** (a) Bimodality of $\log(N_s)$ and first-N3 alignment in CDH-NT1, CDH-NT2, and CDH controls. (b) Sleep-cycle-by-sleep-cycle evolution of the distribution of NREM duration during IRI in NT1 patients, split by presence of a sleep-onset REM period (SOREMP) at the beginning of the night (first row, SOREMP+; second row, SOREMP-). (c) Bimodality of $\log(N_s)$ and first-N3 alignment in SSC.




Figure 14: **Comprehensive transient-oscillation scan in NT1.** Cohen's d (NT1–control, residualized for age, sex, BMI, and site) for all six TO features (rate, prominence, duration, SO-phase coupling strength, SO-phase angle, SO-power); each facet shows six channels (rows) $\times$ five stages $\times$ five frequency bands (columns; band order θ , low- α , slow- σ , fast- σ , β). Dots mark FDR $q < 0.05$; coloured outlines mark the nine direction-based patterns P1–P9 (legend), each defined by a consistent effect sign within a band/stage scope. Hooked arrows ($\hookrightarrow d-i$) on the bottom row of the relevant facets flag the feature shown in topographic detail in the corresponding panel of Figure 6. Circular SO-phase is the raw circular mean (not OLS-residualizable).