

Long-Term Temporal Segmentation and Genotype-Aware Behavioral Phenotyping for AI-Driven Analysis of Sleep–Metabolism Interactions

Manish Dandon

Department of Computer Science, Binghamton University, Binghamton, NY, USA.
manish642@binghamton.edu

Mikkel A. Garrett

School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.
mikkel725@oregonstate.edu

Martins R. Wagner

Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.
martinsrwagner837@ku.edu

Abstract

The bidirectional interplay between sleep architecture and metabolic homeostasis is increasingly recognized as a critical axis of human health, yet its systematic analysis at scale remains limited by methodological fragmentation. This paper presents a systems-level framework that integrates long-term temporal segmentation of multimodal behavioral signals with genotype-aware phenotyping, leveraging contemporary artificial intelligence to dissect sleep–metabolism interactions. We examine the structural trade-offs inherent in deploying memory-efficient video object segmentation models for continuous behavioral state parsing across weeks or months, and we discuss the fusion of such temporally resolved outputs with population-scale genomic data, including the influence of single-nucleotide polymorphisms on metabolic and sleep traits. The discussion centers on architectural design, governance, infrastructure, robustness, fairness, and sustainability rather than on specific algorithmic formulations. We argue that a holistic socio-technical approach is required to reconcile the technical challenges of long-sequence modeling with the ethical imperatives of genetic data stewardship, cross-population fairness, and environmental impact. Policy implications for clinical translation, consumer wellness platforms, and public health monitoring are explored, emphasizing the need for transparent, auditable systems that resist overdetermination by biased training distributions. Through a synthesis of recent advances in computer vision, genomics, and sleep science, we offer a forward-looking perspective on building resilient, equitable, and environmentally conscious AI systems for behavioral phenotyping.

Keywords

temporal segmentation, genotype-aware phenotyping, sleep, metabolism, artificial intelligence, system architecture, data governance, fairness, sustainability.

1. Introduction

Sleep and metabolism are deeply intertwined physiological domains whose dysregulation underlies a wide spectrum of non-communicable diseases, from type 2 diabetes to cardiovascular pathology. Observational and experimental studies have demonstrated that even modest sleep curtailment can perturb glucose homeostasis, alter appetite-regulating hormones, and shift substrate oxidation patterns [1]. Conversely, metabolic states, including hormonal fluctuations and nutrient-sensing pathways, feed back to modulate sleep architecture and circadian timing [2]. Despite these insights, the mechanistic coupling remains poorly resolved in ecologically valid, free-living settings where behavioral rhythms unfold over extended periods and are shaped by both genetic predisposition and environmental contingencies. Traditional analytic approaches in sleep medicine and metabolism have relied on laboratory polysomnography or short-term actigraphy, which offer limited temporal breadth, and on cross-sectional genetic association studies that treat sleep and metabolic traits as static phenotypes rather than dynamic, time-varying constructs.

Recent advances in artificial intelligence, particularly in long-sequence video object segmentation and in the integration of genotype data with deep phenotyping, present an opportunity to reimagine how sleep–metabolism interactions are studied. On the one hand, memory-efficient fine-tuning strategies for foundational segmentation models, such as the Segment Anything Model 2 (SAM2), now permit the parsing of behavioral states from continuous video and wearable sensor streams over weeks and months without exhausting computational memory budgets [3]. On the other hand, large-scale biobank resources have enabled the systematic annotation of genetic variation that moderates changes in gene expression and splicing in response to lifestyle interventions such as exercise and diet-induced weight loss, foreshadowing a future in which genotype-aware metabolic phenotyping is routine [4]. The convergence of these two streams—temporally granular behavioral segmentation and personalized genomic context—creates a powerful platform for AI-driven investigation of sleep–metabolism dynamics.

However, the realization of such a platform is not merely a technical integration problem. It demands a rigorous examination of system architecture, data governance, fairness, sustainability, and deployment trade-offs. This paper addresses these dimensions from a systems science perspective, emphasizing the structural and institutional choices that shape whether these technologies serve public health equitably or amplify existing disparities. We do not focus on specific algorithmic novelty but rather on the design of the socio-technical infrastructure necessary to support robust, long-term behavioral phenotyping in heterogeneous populations. In doing so, we aim to bridge the often-disconnected conversations in computer vision, genomics, sleep medicine, and science and technology policy.

2. Background and Related Work

The epidemiological and physiological evidence linking sleep duration and quality to metabolic outcomes is robust. For instance, early controlled studies established that acute sleep restriction reduces insulin sensitivity and increases evening cortisol levels [1]. More recent work has documented that polymorphisms in circadian clock genes and metabolic regulators jointly influence an individual's vulnerability to sleep-loss-induced metabolic impairment, indicating a genotype–environment interaction that conventional single-modality studies fail to capture [5]. At the cellular level, novel genetically encoded sensors have illuminated that ionic stress, specifically proton gradients, can function as a sleep driver, providing a mechanistic bridge between metabolic byproducts and sleep-regulating neuronal

circuits [6]. This molecular insight reinforces the need for analytical frameworks that span multiple scales, from ion channels to organismal behavior.

On the computational side, the segmentation of behavioral states from sensor data has progressed dramatically. The original Segment Anything model introduced a paradigm for general-purpose image segmentation, and its video extension, SAM2, applied a streaming memory architecture to maintain context across arbitrarily long sequences [3]. Nonetheless, applying such models to continuous, weeks-long behavioral recordings from in-home environments presents a distinct set of challenges: illumination variation, occlusions, diverse subject postures, and storage constraints. Memory-aware fine-tuning methods have been developed that adapt the self-attention mechanisms of SAM2 to significantly reduce memory footprint while preserving temporal coherence, making it feasible to process month-scale video without catastrophic forgetting or memory overflow [7]. These advances in object segmentation provide the low-level behavioral primitives—sleep postures, eating episodes, ambulation—that can be strung together into higher-order behavioral phenotypes.

Parallel developments in genomics have produced reference datasets and analytic methods for genotype-aware phenotyping. The UK Biobank and similar cohorts have shown that deep behavioral and imaging phenotypes, when combined with genome-wide genotyping, can reveal novel loci for complex traits [8]. Recent work on exercise and diet-induced weight loss in human skeletal muscle has demonstrated how polymorphisms affect not only baseline gene expression but also the magnitude and direction of splicing changes in response to metabolic stimulus, underscoring the dynamic nature of genetic influence [4]. These dynamics are precisely the kind of temporal interaction that long-term behavioral segmentation could illuminate: an individual's genetic architecture may moderate not only average sleep need but also the day-to-day variability in sleep–metabolism coupling.

However, integrating these disparate data types is not straightforward. Existing efforts have typically analyzed sleep and metabolic endpoints in aggregate, discarding the rich temporal structure that segment-level behavioral data provide. Conversely, computer vision pipelines have largely treated sleep stage classification or activity recognition as standalone tasks without linking to genetic or metabolic covariates. The present work argues that a purposeful systems design can overcome this fragmentation.

3. System Architecture for Long-Term Temporal Segmentation

Designing an AI system capable of segmenting behavioral states across months of continuous recording requires navigating a tension between temporal fidelity and computational tractability. A naive application of transformer-based video segmentation models to year-long sequences would entail a memory complexity that is quadratic in sequence length, quickly exceeding the capacity of even high-performance computing clusters. Memory-aware fine-tuning, such as that proposed in MFT for SAM2, addresses this by introducing sparse memory tokens and hierarchical temporal abstractions, reducing the effective context window to manageable chunks without discarding long-range dependencies [7]. In a deployment context, this implies a modular architecture where a streaming ingestion layer buffers raw sensor feeds, a segmentation core applies the fine-tuned model to produce frame-level behavioral labels, and a downstream aggregation module compiles these labels into occurrence metrics and circadian statistics.

The choice of sensor modalities introduces further architectural trade-offs. Video provides the highest information density but raises privacy concerns and imposes data transmission

burdens. Wearable inertial sensors and under-mattress ballistocardiography offer less intrusive alternatives, yet they yield coarser behavioral segmentation and may fail during sensor displacement or low-adherence periods. A hybrid architecture that dynamically weights modalities based on confidence scores and availability can offer resilience, but it demands careful calibration to avoid compounding errors when primary modalities are degraded. The system must also incorporate an ontology of behavioral states that is both clinically meaningful and learnable from sensor data; sleep–wake classification alone is insufficient when the goal is to link metabolic outcomes to specific sleep postures, nocturnal eating, or early-morning physical activity.

Storage and computational efficiency are critical for sustainability and scalability. Long-term behavioral monitoring generates terabytes of data per subject per year. In-memory processing of such volumes is infeasible without aggressive compression and on-the-fly feature extraction. Architectures that embed lightweight edge processes for initial segmentation, with periodic cloud synchronization for genomic integration and model updating, can balance responsiveness with resource consumption. However, this edge-cloud split introduces synchronization and consistency challenges, particularly when genomic data are protected by strict governance policies that forbid their transmission to edge devices. A tiered trust model, wherein only anonymized behavioral features are shared upstream while raw genetic data remain within a secure enclave, becomes a necessary architectural component.

4. Genotype-Aware Behavioral Phenotyping: Integrating Genetic Variation

The integration of genetic data into behavioral phenotyping pipelines enables the personalization of sleep–metabolism analysis but also exposes deep structural tensions around privacy, equity, and interpretability. Genome-wide association studies have identified hundreds of loci associated with sleep duration, chronotype, and insomnia, and metabolic traits such as body mass index and insulin resistance are similarly polygenic [5]. When these genetic scores are combined with temporally resolved behavioral sequences, one can interrogate how genetic risk factors manifest through specific behavioral patterns over time—for example, whether carriers of certain variants exhibit delayed postprandial activity or fragmented sleep after high-fat meals.

The recent demonstration that polymorphisms moderate splicing changes during exercise and dietary interventions exemplifies the dynamic nature of genetic influence that long-term segmentation can capture [4]. If a system detects that an individual’s metabolic response to sleep loss varies dramatically across days, a genotype-aware model could attribute part of this variance to the interaction between circadian gene variants and day-specific behavioral context. However, the use of polygenic scores raises well-documented concerns about portability across ancestries: scores derived predominantly from European-ancestry cohorts often lose predictive power in other populations, potentially leading to mischaracterization of risk and inappropriate behavioral recommendations [9]. Therefore, any deployed system must incorporate methods for ancestry-aware calibration and uncertainty quantification, rather than outputting point estimates of genetic liability.

Furthermore, genotype data are inherently identifiable and subject to stringent regulatory frameworks such as the General Data Protection Regulation and the Genetic Information Nondiscrimination Act. Architecturally, this necessitates a separation of genomic processing from the behavioral segmentation front-end, possibly via a federated learning scheme where genetic models are updated using sanitized behavioral summaries without exposing raw DNA sequences. The governance implications of storing and processing genomic data alongside

continuous behavioral streams are profound: data breaches could simultaneously reveal a person's location patterns, sleep habits, and genetic vulnerabilities, creating an unprecedented privacy risk profile [10]. System design must therefore incorporate end-to-end encryption, differential privacy during model training, and a clear data lifecycle policy that specifies retention limits and user-consent revocation mechanisms.

5. Data Infrastructure, Privacy, and Governance

The infrastructure required to support AI-driven sleep–metabolism analysis extends beyond computational hardware to include institutional processes for data stewardship, consent management, and cross-jurisdictional compliance. Longitudinal behavioral monitoring inevitably collects information about intimate aspects of daily life, including sleep arrangements, intimate partner interactions, and timing of substance use. Data subjects must be able to provide granular, dynamic consent that can be withdrawn for specific data streams without invalidating the entire longitudinal record, and such withdrawal must be technically enforceable, not merely aspirational. This challenges standard data lake architectures, which often assume immutable append-only logs.

Federated learning has been proposed as a privacy-preserving alternative to centralized training, yet its application to multimodal, temporally structured data remains immature. Synchronizing model updates across heterogeneous edge nodes with varying data distributions risks introducing statistical heterogeneity that degrades global model performance and fairness [11]. Moreover, federated learning alone does not prevent inference of sensitive attributes from model updates; differential privacy guarantees must be layered on top with careful accounting of the privacy budget over time, given the continuously accruing nature of behavioral data. The integration of genetically encoded biosensors, such as ionic-stress reporters that link proton dynamics to sleep drive [6], into wearable devices could eventually generate molecular data streams that further complicate this privacy calculus, as they would constitute both behavioral and physiological data with potential diagnostic implications.

Governance frameworks must also address the secondary use of behavioral–genomic datasets. Partnerships with consumer device manufacturers and biobanks create value but also introduce conflicts of interest and data commodification risks. A system designed for long-term sleep–metabolism research risks repurposing for insurance risk stratification or workplace surveillance unless transparent governance boundaries and independent audit mechanisms are established at the design stage. Institutional review boards and data ethics committees must be equipped to evaluate the unique risks of temporally dense, genetically linked behavioral data, which fall outside the purview of conventional static biobank regulations [12]. A model of data trusts, wherein data subjects retain collective governance over data use, may offer a path toward sustainable, ethical data infrastructure.

6. Robustness, Fairness, and Bias Mitigation

AI systems for behavioral phenotyping are vulnerable to multiple forms of bias that interact in pernicious ways. Sensor-related biases arise because the performance of video-based sleep segmentation drops under low-light conditions, which disproportionately affects lower-income households that lack ambient night lighting, while wearable-based systems may underperform on darker skin tones due to photoplethysmography signal attenuation [13]. Genomic biases stem from the overrepresentation of European-ancestry populations in reference panels, which limits the accuracy of genotype-aware predictions in non-European

groups and may lead to the development of interventions that are less effective or even harmful for those groups [9]. Temporal biases occur when models trained predominantly on weekday sleep patterns fail to characterize weekends or shift-work schedules, marginalizing populations whose occupational rhythms deviate from the norm.

Addressing these biases requires a system-level strategy, not merely algorithmic debiasing post hoc. During data collection, stratified sampling protocols that oversample underrepresented demographic and circadian groups can improve the representativeness of the training distribution. At the model level, multi-task learning with fairness constraints can encourage predictive parity across demographic groups without sacrificing overall accuracy, though such constraints must be chosen carefully to avoid inadvertently penalizing groups with genuinely different phenotype distributions [14]. Continuous monitoring of model outputs with demographic disaggregation is essential, yet it raises a tension with privacy, as it requires collecting protected attributes that individuals may be reluctant to disclose. Privacy-preserving fairness auditing techniques, such as those based on secure multiparty computation, remain an active area of development but are not yet robust enough for production behavioral systems.

Robustness to distributional shift is equally critical. A segmentation model trained in one geographic region or season may fail when deployed in different climates, housing types, or cultural sleep practices. The inclusion of environmental light sensors and temperature data can provide covariates that partially absorb these domain shifts, but a more principled approach is to maintain a continuous update cycle with human-in-the-loop verification that prevents model drift from compounding unchecked. This need for ongoing supervision raises the question of who bears the cost of model maintenance over years of operation, a challenge often overlooked in time-limited research grants but central to long-term sustainability.

7. Deployment, Scalability, and Sustainability

Translating an AI-driven sleep–metabolism analysis platform from research prototype to population-scale deployment introduces logistical, economic, and environmental considerations. A system that must process high-resolution video from millions of users would consume enormous amounts of energy, contributing to the carbon footprint of digital health infrastructure. The use of memory-aware fine-tuning reduces computational demand during inference by compressing temporal context, but training such models on diverse, longitudinally annotated datasets remains energy-intensive [7]. A full lifecycle environmental assessment, from data center construction through end-of-device disposal, should inform system architecture decisions, aligning with the growing call for sustainable AI practices [15].

Scalability is not solely a computational issue; it also encompasses clinical and social scalability. The behavioral phenotyping outputs must be interpretable to healthcare providers and patients without requiring specialist data science knowledge. This demands a presentation layer that translates segment-level predictions into clinically actionable summaries—for instance, a moving average of sleep fragmentation linked to post-dinner glucose excursions—while conveying the uncertainty inherent in automated measurements. Furthermore, the system must integrate with electronic health records and existing digital health ecosystems, which are often fragmented across proprietary standards. Open API specifications and interoperability mandates may be necessary to prevent vendor lock-in and to enable the kind of multi-institutional data sharing that genetic discovery requires.

Economic sustainability hinges on a value proposition that aligns the interests of users, healthcare payers, and device manufacturers. If the system demonstrates that early detection of sleep–metabolism dysregulation can reduce the incidence of type 2 diabetes or cardiovascular events, payers may be willing to subsidize deployment. However, the risk of exaggerated claims based on short-term validation studies is high, and rigorous longitudinal trials in diverse cohorts must precede widespread clinical endorsement. In parallel, a consumer-facing tier could provide general wellness insights without medical claims, but this blurs the boundary between lifestyle advice and medical device regulation, a domain where policy is rapidly evolving.

8. Policy Implications and Future Directions

The widespread deployment of AI systems that integrate long-term behavioral segmentation with genotype-aware phenotyping will require updated regulatory frameworks. Current medical device regulations, such as the FDA’s 510(k) pathway and the EU Medical Device Regulation, were not designed for adaptive AI models that update continuously from streaming data. The concept of predetermined change control plans, wherein manufacturers pre-specify the boundaries of permissible model evolution, represents a promising regulatory innovation, but its application to systems that combine computer vision, genomics, and metabolic sensing remains untested. Regulators will need multidisciplinary expertise to evaluate whether these combined systems amplify biases in ways that no single component would in isolation.

Policy must also address the societal implications of making genotype-aware sleep–metabolism insights widely accessible. There is a risk that such information could reinforce a narrow, genetic determinism that overlooks structural determinants of health, such as food deserts, shift-work labor practices, and noise pollution, which profoundly influence both sleep and metabolism. A responsible deployment framework would explicitly model environmental and social covariates and would be transparent about the limits of genetic prediction. Public education campaigns will be necessary to convey that a “high-risk” polygenic score does not inevitably lead to metabolic disease and that behavioral context, captured through segmentation, mediates genetic risk.

Looking forward, the incorporation of novel molecular sensors into long-term behavioral monitoring could dramatically deepen phenotype resolution. For example, if wearable or implantable devices could detect ionic-stress dynamics or proton gradients associated with sleep pressure [6], the temporal segmentation system would gain a direct biochemical layer that complements behavioral observations. This would further blur the boundary between research tool and medical device, intensifying regulatory scrutiny. Additionally, the emerging field of memory-aware fine-tuning for video segmentation suggests that future systems could operate entirely on-device, obviating the need for cloud transmission of sensitive behavioral data and strengthening privacy guarantees [7]. However, on-device processing shifts the computational carbon load to millions of battery-powered devices, creating new sustainability challenges that must be managed through efficient model compression and renewable-energy charging ecosystems.

9. Conclusion

The convergence of long-term temporal segmentation and genotype-aware behavioral phenotyping offers an unprecedented opportunity to unravel the complex dynamics of sleep and metabolism in large, diverse populations. Yet the path from technical feasibility to

equitable, sustainable impact is fraught with structural challenges that transcend algorithmic optimization. This paper has argued that the design of AI systems for sleep–metabolism analysis must be guided by a holistic systems perspective that addresses architectural trade-offs, privacy-preserving data governance, fairness across demographic groups, and the lifecycle environmental costs of computation. By anchoring the discussion in concrete advances—memory-aware video segmentation, the role of polymorphisms in metabolic responses, and novel ionic-stress sensors—we have shown that these challenges are both urgent and addressable. Future work must engage not only computer scientists and geneticists but also clinicians, ethicists, regulators, and the communities whose sleep and metabolic health stand to benefit or be harmed by these technologies.

References

1. Spiegel, K., Leproult, R., & Van Cauter, E. (1999). Impact of sleep debt on metabolic and endocrine function. *The Lancet*, 354(9188), 1435–1439.
2. de Zambotti, M., Cellini, N., Goldstone, A., Colrain, I. M., & Baker, F. C. (2019). Wearable sleep technology in clinical and research settings. *Medicine & Science in Sports & Exercise*, 51(7), 1538–1557.
3. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., ... & Girshick, R. (2023). Segment anything. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.
4. Wang, W. et al. Impact of polymorphisms on gene expression and splicing in response to exercise and diet-induced weight loss in human skeletal muscle tissues. *Cell Genom* 5, 100951 (2025).
5. Lane, J. M., Liang, J., Vlasac, I., Anderson, S. G., Bechtold, D. A., Bowden, J., ... & Saxena, R. (2017). Genome-wide association analyses of sleep disturbance traits identify new loci and highlight shared genetics with neuropsychiatric and metabolic traits. *Nature Genetics*, 49(2), 274–281.
6. Ji, Z., Liu, J., Wang, B., Wei, S., Bian, Y., Zeng, W., Chung, C. I., Ma, Z., Zhang, J., Shu, X., & Ma, D. K. (2026). A genetically encoded ionic-stress sensor reveals protons as a sleep driver. *bioRxiv*. <https://doi.org/10.64898/2026.01.27.702131>
7. Li, G., Yuan, H., Chen, S., Hu, Q., Wang, J., & Jiang, K. (2026). MFT: Memory-Aware Fine-Tuning of SAM2 for Efficient Long-Sequence Video Object Segmentation. *IEEE Signal Processing Letters*.
8. Bycroft, C., Freeman, C., Petkova, D., Band, G., Elliott, L. T., Sharp, K., ... & Marchini, J. (2018). The UK Biobank resource with deep phenotyping and genomic data. *Nature*, 562(7726), 203–209.
9. Martin, A. R., Kanai, M., Kamatani, Y., Okada, Y., Neale, B. M., & Daly, M. J. (2019). Clinical use of current polygenic risk scores may exacerbate health disparities. *Nature Genetics*, 51(4), 584–591.
10. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 2053951716679679.
11. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2), 1–210.

12. Floridi, L., Cows, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., ... & Schafer, B. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707.
13. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453.
14. Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. *New England Journal of Medicine*, 380(14), 1347–1358.
15. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645–3650.
16. Nedeltcheva, A. V., Kilkus, J. M., Imperial, J., Kasza, K., & Penev, P. D. (2010). Sleep curtailment is accompanied by increased intake of calories from snacks. *The American Journal of Clinical Nutrition*, 89(1), 126–133.
17. Reutrakul, S., & Van Cauter, E. (2018). Sleep influences on obesity, insulin resistance, and risk of type 2 diabetes. *Metabolism*, 84, 56–66.
18. Perez-Pozuelo, I., Zhai, B., Palotti, J., Mall, R., Aupetit, M., Garcia-Gancedo, C., ... & Fernandez-Luque, L. (2020). The future of sleep health: a data-driven revolution in sleep science and medicine. *NPJ Digital Medicine*, 3(1), 42.
19. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318.
20. Lea, C., Flynn, M. D., Vidal, R., Reiter, A., & Hager, G. D. (2017). Temporal convolutional networks for action segmentation and detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 156–165.
21. Tam, V., Patel, N., Turcotte, M., Bossé, Y., Paré, G., & Meyre, D. (2019). Benefits and limitations of genome-wide association studies. *Nature Reviews Genetics*, 20(8), 467–484.
22. Depner, C. M., Stothard, E. R., & Wright, K. P. (2014). Metabolic consequences of sleep and circadian disorders. *Current Diabetes Reports*, 14(7), 507.