

# Graph Neural Networks for Multisensor Physiological State Modeling in Human Activity Monitoring

Nils Perkins

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.

nilswork@uab.edu

Clifford M. Burns

Department of Computer Science, Binghamton University, Binghamton, NY, USA.

clifford.m.burns@binghamton.edu

Nils Steele

Department of Computer Science, University of Central Florida, Orlando, FL, USA.

nilssteele79@ucf.edu

Dustin Morales

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.

dustin.work@buffalo.edu

## Abstract

Human activity monitoring has expanded from classifying locomotion to estimating complex physiological states from multiple wearable sensors. These sensor streams are heterogeneous, asynchronous, and physiologically coupled, which makes flat feature representations insufficient for reliable inference. Graph neural networks offer a principled relational inductive bias for modeling interacting physiological subsystems. This paper presents a systems-oriented analysis of graph neural networks for multisensor physiological state modeling in human activity monitoring. It examines the shift from Euclidean and sequence-based sensor processing to graph-based relational inference, the architectural choices that shape graph construction and message passing, and the deployment trade-offs across cloud, edge, and on-device settings. The discussion further addresses data infrastructure requirements, robustness under distribution shift, uncertainty quantification, and regulatory alignment. Fairness and governance are treated as central design constraints because physiological sensing can encode population-specific differences and sensitive health information. The paper also considers sustainability, long-term model evolution, and institutional accountability. By integrating these dimensions, the analysis moves beyond predictive performance to clarify the structural trade-offs that determine whether graph-based physiological state modeling can be deployed responsibly in wearable, clinical, and occupational environments.

## Keywords

graph neural networks; multisensor fusion; physiological state modeling; human activity monitoring; wearable systems; responsible deployment.

## 1. Introduction

Human activity monitoring has evolved from classifying simple locomotion patterns to estimating fine-grained physiological states such as cardiovascular load, respiratory effort, thermoregulatory stress, and autonomic arousal. Wearable devices now integrate photoplethysmography, electrodermal activity, skin temperature, accelerometry, and sometimes electrocardiography, producing heterogeneous streams that are temporally misaligned and context dependent. A central problem in modern multisensor analytics is that physiological subsystems are not independent channels. They interact through autonomic regulation, mechanical coupling, vascular dynamics, and behavioral context. Treating the sensor suite as a flat vector or a fixed grid therefore discards structural information that is essential for robust inference. Graph neural networks provide a formalism for learning over irregular relational domains and have matured rapidly as a class of deep models [1]. Early formulations such as graph convolutional networks [2], attention-based graph layers [3], and inductive embedding approaches [4] demonstrated that relational inductive biases can improve learning efficiency and generalization. Comprehensive surveys have established the breadth of graph-based learning across domains [5]. In parallel, wearable activity recognition has built a substantial empirical foundation around inertial and multimodal sensing [6]. However, the intersection of graph-based learning and multisensor physiological state modeling remains underexplored from a systems perspective. This paper examines that intersection with attention to architectural trade-offs, data infrastructure, robustness, fairness, governance, and long-term sustainability.

The motivation for a systems-level perspective is that physiological state modeling is not only a machine learning problem. It touches on sensor reliability, user consent, clinical validity, energy consumption, and institutional oversight. A graph neural architecture may improve sensor fusion, but its value depends on whether it can operate reliably in real-world conditions and whether its errors are legible to clinicians, regulators, and users. The following sections therefore analyze the relational nature of physiological sensing, the structural choices that shape graph-based models, the infrastructure needed for deployment, the robustness and regulatory constraints that define acceptable operation, and the broader sociotechnical implications of inferring physiological states from multisensor data.

## **2. Multisensor Physiological State Modeling as a Relational Inference Problem**

Traditional activity monitoring has relied heavily on body-worn inertial sensors and benchmark datasets that support comparison across recognition pipelines [7]. Physiological state modeling extends the problem in several ways. Signals such as photoplethysmography are highly susceptible to motion artifacts, and recovering the underlying waveform during physical activity requires specialized conditioning, including adaptive notch-filtration [8]. Once cleaned, physiological channels still exhibit complex coupling. Deep convolutional and recurrent models have shown strong results for multimodal wearable activity recognition by learning temporal features from raw sensor streams [9][10]. However, these models typically assume that sensor channels are aligned along a shared time axis and do not explicitly represent pairwise relationships among physiological variables.

A relational view is more natural because physiological signals arise from interacting regulatory systems. A graph can represent each sensor channel, derived physiological variable, body location, or contextual state as a node. Edges encode anatomical proximity, shared regulatory pathways, or statistical dependence. Message passing then allows a model to update each node representation using information from functionally related nodes. This is different from sensor fusion that concatenates features or merges modalities through attention

over time. The graph structure can be partially known from physiology and partially learned. In a running scenario, accelerometer-derived cadence and photoplethysmographic heart rate share a mechanical coupling. In a stress protocol, electrodermal activity and skin temperature exhibit delayed autonomic coupling. DeepSense demonstrated that unified deep learning over mobile sensing time series can outperform modality-specific processing [11], but its architecture remains grid-like and sequential. Graph neural networks extend this line of work by making the interaction graph explicit and adaptive. Surveys of deep learning for sensor-based activity recognition have noted the need for architectures that can capture dependencies beyond simple temporal convolution [12].

Several system-level challenges reinforce the value of relational modeling. Sensor channels are often asynchronous, with different sampling rates and latencies. Some channels may be missing for intervals because of contact loss, battery depletion, or user removal of the device. Calibration drift can change signal distributions across days. A graph representation does not automatically solve these problems, but it provides a structural mechanism for handling partial observations. When one node is unavailable, message passing can continue over the remaining graph, allowing inference to degrade gracefully rather than fail entirely. The graph can also incorporate contextual nodes such as activity type, time of day, or ambient temperature, enabling the model to condition physiological interpretation on external circumstances.

### **3. Architectural Considerations for Graph-Based Physiological Representation**

Designing a graph neural architecture for physiological state modeling begins with graph construction, which is often more consequential than layer choice. The graph may be static, with nodes arranged according to body locations and known physiological connections, or dynamic, with edges reweighted according to activity context. Static anatomical graphs are interpretable and inexpensive, but they may miss transient couplings such as respiratory sinus arrhythmia under varying load. Fully learned adjacency matrices are flexible but can overfit small physiological datasets. A hybrid approach uses anatomical priors to initialize the graph and allows attention weights to modify edge strengths. Gated graph sequence networks introduce stateful message passing that is useful when physiological dynamics unfold over multiple timescales [13]. Attention-based graph layers [3] and transformer-style architectures [14] offer fine-grained reweighting but require careful regularization to prevent attention from collapsing onto high-magnitude sensor channels.

Multi-task learning is particularly relevant because physiological state variables are mutually informative. A graph model can share latent representations across tasks such as heart rate estimation, respiratory rate inference, and fatigue prediction. Multi-task architectures can improve data efficiency and enforce consistency, but they also introduce negative transfer when tasks conflict or when some outputs have weak labels [15]. The system designer must decide whether to use shared graph encoders with task-specific heads or to separate encoders for groups of related variables. Hierarchical pooling can create a multiscale representation from sensor nodes to body regions to whole-body physiological state. This supports both fine-grained sensor anomaly detection and systemic state estimation. In comparison, non-relational baselines such as gradient-boosted trees remain strong for tabular features and are often easier to deploy [16]. The benefit of a graph model emerges when the relational structure is sufficiently stable and the data volume justifies end-to-end learning.

Architectural choices also interact with labeling constraints. Dense supervision is rare in wearable studies because ground-truth physiological states often require clinical equipment.

Graph networks can be trained with weak or partial labels by using co-training across related nodes and tasks. They can also incorporate inductive biases that allow generalization to unseen sensor configurations, which is critical when devices are upgraded or body placement changes [4]. The system-level trade-off is between model expressiveness, memory footprint, and the reliability of learned edges. Attention-based graph layers increase capacity but may demand more energy and heighten the risk of silent failure when the attention distribution becomes uninformative. These trade-offs must be evaluated within the constraints of the deployment environment rather than purely in benchmark settings.

#### **4. Data Infrastructure and Deployment Topologies**

Deploying graph-based physiological state modeling requires a data infrastructure that can ingest, align, clean, and label multisensor streams at scale. The data lifecycle includes sensor calibration, synchronization, segment labeling, graph construction, model training, validation, and continuous monitoring. Wearable data are often collected in uncontrolled environments, producing missing segments, clock drift, and variable sampling rates. A graph representation can absorb some of this irregularity by allowing nodes to carry asynchronous observations, but the infrastructure must still maintain lineage from raw sensor records to graph snapshots. Benchmarks for graph neural networks have clarified that architectural gains can be masked by inconsistent evaluation, data leakage, or poor splits [17]. Similar discipline is needed for physiological graphs because overlapping subjects or sensor sessions can easily create leakage across train and test sets.

Deployment topologies range from cloud-centric inference to fully on-device execution. Cloud training benefits from large aggregated datasets and high-performance accelerators, but it raises privacy and latency concerns. On-device inference protects sensitive physiological data and enables real-time feedback, yet graph models may be too large for wearable processors. Edge aggregation points, such as a smartphone or a home hub, offer a middle path. The graph can be partitioned across the body-area network and the edge node, with lightweight local encoders on each sensor and a centralized graph processor at the edge. Training large graph networks also introduces scalability challenges. Sampling methods are used to approximate message passing over large dynamic graphs without loading the full neighborhood [18]. These methods reduce memory consumption but can introduce variance that must be controlled in safety-critical applications.

Self-supervised pretraining can reduce dependence on labeled physiological data. Graph-level representation learning via mutual information maximization [19] and redundancy-reduction objectives [20] can leverage unlabeled multisensor records to learn reusable node and edge embeddings. These pretrained representations can then be fine-tuned for downstream tasks such as stress detection, recovery assessment, or clinical deterioration prediction. From an infrastructure standpoint, pretraining requires substantial storage and computational capacity, but it amortizes well across multiple downstream services. Governance mechanisms must track which pretraining corpus was used, how data were processed, and whether the learned graph encodes device-specific artifacts. Without such lineage, a model may inherit biases or calibration assumptions that are not visible at deployment time.

#### **5. Robustness, Uncertainty, and Regulatory Alignment**

Robustness is a first-order concern in physiological state modeling because errors can trigger false alarms or missed interventions. Sensor noise, motion artifacts, skin pigmentation, perfusion variability, and electrode displacement all perturb the input graph. Graph models

can be made more robust by denoising the input signals before graph construction, regularizing edge weights, and training with aggressive data augmentation that simulates missing nodes and noisy edges. Upstream signal recovery remains essential before relational modeling can succeed, as demonstrated by adaptive notch-filtration for photoplethysmographic signals during physical activity [8]. Even with cleaning, models should estimate uncertainty rather than produce point predictions alone. Uncertainty estimation is particularly important when the graph is incomplete or when a sensor is temporarily disconnected.

Regulatory alignment shapes the architecture. If the system is deployed as a wellness tool, the evidentiary burden is lower, but if it informs clinical decisions, it may be regulated as a medical device. Regulatory frameworks generally require documented risk management, validation on representative populations, and traceability of model changes. Graph-based models complicate traceability because learned edges and message passing can obscure the causal chain from input sensor to output state. Explainability methods for graphs are less mature than feature attribution methods for tabular or image models. A manageable approach is to keep a subset of the graph anatomically constrained and to provide auxiliary outputs that flag which nodes and edges most influenced a prediction. This can support clinical review without sacrificing the benefits of learned relational features. Bias and fairness must also be evaluated during validation, not after deployment [21].

Distribution shift is especially pronounced in physiological monitoring. A model trained on healthy young adults may not transfer to older adults, pregnant individuals, or people with cardiovascular conditions. Seasonal temperature, medication, and device placement create additional shifts. Continuous monitoring requires drift detection, periodic recalibration, and possibly federation across institutions. Graph models offer a structural advantage because new sensor modalities or body locations can be added as nodes without redesigning the entire input pipeline. However, this flexibility also demands governance that controls how new edges are authorized and how legacy predictions are interpreted after graph modifications. Such governance is not a secondary concern; it determines whether the system remains safe and defensible over time.

## **6. Fairness, Governance, and Sociotechnical Implications**

Fairness in multisensor physiological monitoring is not merely a matter of balanced datasets. Sensor signals themselves can differ systematically across skin tones, body composition, age, and health status. Photoplethysmography, for example, is known to be affected by melanin concentration and perfusion. If such differences are not explicitly considered, graph models may learn spurious associations between demographic proxies and physiological states. Surveys of bias and fairness in machine learning distinguish among multiple fairness definitions that may be mutually incompatible [21][22]. System designers must therefore choose context-specific fairness objectives, document their limitations, and assess impact across subpopulations.

Governance for physiological state modeling must address consent, secondary use, and data sovereignty. Multisensor records are deeply personal and can reveal sleep, stress, medication adherence, and disease progression. Users may not anticipate that graph-based inference can recover sensitive states from combinations of seemingly innocuous channels. Data governance frameworks should specify the purposes for which models may be used, the retention period for raw signals and graph embeddings, and the conditions under which data may be shared with researchers or insurers. Fairness and machine learning scholarship

emphasizes that fairness cannot be reduced to a single metric and that institutional processes are needed to adjudicate competing values [23].

Policy implications extend to workplace wellness programs, insurance underwriting, and clinical remote monitoring. If graph-based physiological state models are used to infer fatigue or stress in occupational settings, employees may face pressure to wear sensors or accept algorithmic assessments. Transparent opt-in policies, independent audits, and collective bargaining may be necessary to prevent coercive monitoring. At the same time, these systems can improve safety in physically demanding jobs and support early detection of deterioration in chronic disease. A sociotechnical perspective treats the graph model, sensor infrastructure, institutional incentives, and user practices as a single integrated system rather than as a technical artifact with external social effects. This perspective is essential for aligning model capabilities with legitimate social purposes.

## **7. Sustainability and Long-Term System Evolution**

Sustainability encompasses energy consumption, hardware lifecycle, and the maintenance burden of large graph-based models. Training deep graph networks on high-frequency sensor streams can consume substantial energy, especially if experiments require repeated hyperparameter searches. The environmental cost must be weighed against the clinical or operational benefit. In industrial human activity recognition, deep learning models have been widely applied to improve safety and productivity, but the field has also recognized the need for efficient architectures and meaningful performance baselines [24]. For physiological monitoring, sustainability also means designing devices and algorithms that remain useful over years rather than requiring constant replacement.

Long-term evolution poses a different set of challenges. Sensor hardware changes, firmware updates, and shifts in population health alter the data distribution. A graph model that performs well at deployment may degrade silently if the learned edges no longer reflect physiological couplings. Continual learning can update node and edge representations, but it risks catastrophic forgetting and requires robust evaluation protocols. Versioning must extend beyond model weights to include graph schemas, preprocessing pipelines, and calibration parameters. Institutional memory is necessary to understand why a model was updated and which predictions were produced by which version.

Future systems may integrate graph-based physiological state modeling into digital twins, smart homes, and clinical command centers. In these settings, the graph can expand beyond a single body to include environmental sensors, care team actions, and longitudinal outcomes. This expansion offers new capabilities but also increases the stakes for governance. Public health authorities, healthcare providers, technology vendors, and wearable users will need shared standards for data interoperability, model reporting, and adverse event investigation. The long-term value of graph neural networks in this domain will depend not only on predictive accuracy but also on whether they can be deployed in ways that are explainable, equitable, and institutionally accountable.

## **8. Conclusion**

Graph neural networks offer a compelling framework for multisensor physiological state modeling in human activity monitoring because they align naturally with the relational structure of physiological systems. This paper has argued that the benefits of graph-based modeling must be assessed through a systems lens that includes architectural design, data infrastructure, deployment topology, robustness, regulatory alignment, fairness, governance,

and sustainability. The graph formulation allows heterogeneous sensor channels and physiological variables to be represented as interacting nodes, enabling more flexible fusion and graceful degradation under missing data. However, this flexibility introduces new responsibilities. Graph construction must be disciplined to avoid overfitting, learned edges must be interpretable enough for clinical and regulatory review, and the data lifecycle must preserve provenance from raw signals to graph snapshots.

Deployment choices have structural consequences. Cloud training, edge aggregation, and on-device inference each shift the balance among privacy, latency, energy, and model capacity. Robustness and uncertainty quantification are critical because physiological errors can carry direct human costs. Fairness cannot be assumed from architectural choice alone; it requires deliberate evaluation across populations and transparent governance. Long-term sustainability depends not only on computational efficiency but also on the institutional capacity to maintain, update, and audit models as devices and populations change. The future of graph-based physiological state modeling will be shaped by how well researchers, engineers, clinicians, regulators, and users coordinate around these interdependent concerns.

## References

1. Bronstein, M. M., Bruna, J., LeCun, Y., Szlam, A., & Vandergheynst, P. (2017). Geometric deep learning: Going beyond Euclidean data. *IEEE Signal Processing Magazine*, 34(4), 18–42. <https://doi.org/10.1109/MSP.2017.2693418>
2. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. *arXiv*. <https://arxiv.org/abs/1609.02907>
3. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph attention networks. *arXiv*. <https://arxiv.org/abs/1710.10903>
4. Hamilton, W. L., Ying, R., & Leskovec, J. (2017). Inductive representation learning on large graphs. *arXiv*. <https://arxiv.org/abs/1706.02216>
5. Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2021). A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1), 4–24. <https://doi.org/10.1109/TNNLS.2020.2978386>
6. Bulling, A., Blanke, U., & Schiele, B. (2014). A tutorial on human activity recognition using body-worn inertial sensors. *ACM Computing Surveys*, 46(3), Article 33. <https://doi.org/10.1145/2499621>
7. Reiss, A., & Stricker, D. (2012). Introducing a new benchmarked dataset for activity monitoring. In *Proceedings of the 16th International Symposium on Wearable Computers* (pp. 108–109). IEEE. <https://doi.org/10.1109/ISWC.2012.13>
8. Zheng, X., Dwyer, V. M., Barrett, L. A., Derakhshani, M., & Hu, S. (2022). Adaptive notch-filtration to effectively recover photoplethysmographic signals during physical activity. *Biomedical Signal Processing and Control*, 72, 103303.
9. Ordóñez, F. J., & Roggen, D. (2016). Deep convolutional and LSTM recurrent neural networks for multimodal wearable activity recognition. *Sensors*, 16(1), 115. <https://doi.org/10.3390/s16010115>
10. Radu, V., Tong, C., Bhattacharya, S., Lane, N. D., Mascolo, C., Marina, M. K., & Kawsar, F. (2018). Multimodal deep learning for activity and context recognition. *Proceedings of*

the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, 1(4), Article 157. <https://doi.org/10.1145/3161174>

11. Yao, S., Hu, S., Zhao, Y., Zhang, A., & Abdelzaher, T. (2017). Yao, S., Hu, S., Zhao, Y., Zhang, A., & Abdelzaher, T. (2017). DeepSense: A unified deep learning framework for time-series mobile sensing data processing. In Proceedings of the 26th International Conference on World Wide Web (pp. 351–360). International World Wide Web Conferences Steering Committee. <https://doi.org/10.1145/3038912.3052577>
12. Wang, J., Chen, Y., Hao, S., Peng, X., & Hu, L. (2019). Deep learning for sensor-based activity recognition: A survey. Pattern Recognition Letters, 119, 3–11. <https://doi.org/10.1016/j.patrec.2018.02.010>
13. Li, Y., Tarlow, D., Brockschmidt, M., & Zemel, R. (2016). Gated graph sequence neural networks. arXiv. <https://arxiv.org/abs/1511.05493>
14. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In Advances in Neural Information Processing Systems 30. Curran Associates.
15. Zhang, Y., & Yang, Q. (2022). A survey on multi-task learning. IEEE Transactions on Knowledge and Data Engineering, 34(12), 5586–5609. <https://doi.org/10.1109/TKDE.2021.3070203>
16. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 785–794). ACM. <https://doi.org/10.1145/2939672.2939785>
17. Dwivedi, V. P., Joshi, C. K., Luu, A. T., Laurent, T., Bengio, Y., & Bresson, X. (2023). Benchmarking graph neural networks. Journal of Machine Learning Research, 24(43), 1–48.
18. Chen, J., Ma, T., & Xiao, C. (2018). FastGCN: Fast learning with graph convolutional networks via importance sampling. arXiv. <https://arxiv.org/abs/1801.10247>
19. Sun, F.-Y., Hoffmann, J., Verma, V., & Tang, J. (2020). InfoGraph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization. arXiv. <https://arxiv.org/abs/1908.01000>
20. Bielak, P., Kajdanowicz, T., & Chawla, N. V. (2022). Graph Barlow Twins: A self-supervised representation learning framework for graphs. Knowledge-Based Systems, 256, 109631. <https://doi.org/10.1016/j.knosys.2022.109631>
21. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. ACM Computing Surveys, 54(6), Article 115. <https://doi.org/10.1145/3457607>
22. Kleinberg, J., Mullainathan, S., & Raghavan, M. (2017). Inherent trade-offs in the fair determination of risk scores. In Proceedings of the 8th Innovations in Theoretical Computer Science Conference (Article 43). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. <https://doi.org/10.4230/LIPIcs.ITCS.2017.43>
23. Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. In Proceedings of the Conference on

Fairness, Accountability, and Transparency (pp. 59–68). ACM.  
<https://doi.org/10.1145/3287560.3287598>

24. Chen, K., Zhang, D., Yao, L., Guo, B., Yu, Z., & Liu, Y. (2021). Deep learning for sensor-based human activity recognition: Overview, challenges, and opportunities. *ACM Computing Surveys*, 54(4), Article 77. <https://doi.org/10.1145/3447744>