

AI-Driven Diagnostics: A Multimodal Framework for Monitoring the Comorbid Cluster of Chronic Pain, Depression, and Sleep Disorders

Tejas Padliya

Westcliff University

ORCID: [0009-0009-7048-6134](https://orcid.org/0009-0009-7048-6134)

Abstract

Chronic pain, depression, and sleep disorders are significant health concerns worldwide because they are often comorbid and reinforce one another through a cycle of suffering. Each of these conditions has a negative impact on quality of life and complicates the delivery of health care services; furthermore, they carry a heavy financial burden. Currently, health care facilities provide care for these conditions separately, on an episodic and reactive basis without accounting for the interactions that take place dynamically. Therefore, we propose the FUSE framework (Fidelity, Uncertainty, Synchronization, Explainability) as four core design principles for trustworthy multimodal AI diagnostic systems and describe the Temporal Cross-Modal Alignment Graph (TCAG) as a way of architecting a tangible implementation of the FUSE framework. We argue that the TCAG architecture addresses three primary barriers to the working application of AI in the digital health space: data heterogeneity, model opacity, and asynchronous sensor streams. We describe how the TCAG architecture integrates time-series data from wearable technologies, smartphones, and electronic health records (EHR) by aligning asynchronous time-series data. In the alignment phase, a graph is produced. The graph is processed through a graph attention mechanism that propagates information across aligned modality nodes and explainable clinical narratives. A structured, reasoned architectural comparison of TCAG against Early and Late Fusion baselines is presented across four dimensions relevant to the clinical deployment of AI: (1) missing data handling, (2) capturing cross-modal interactions, (3) temporal alignment fidelity, and (4) quantifying uncertainty. We further report preliminary empirical support from a cross-sectional analysis of NHANES survey data and a controlled synthetic evaluation in which TCAG showed higher mean discrimination and greater robustness to sensor dropout and clock misalignment. In summary, we offer clinicians an open-loop clinical decision support tool to assist them in treating some of the most complex chronic diseases.

Keywords: digital health, digital biomarkers, chronic pain, depression, sleep disorders, multimodal data fusion

Licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.

Introduction

Managing chronic conditions, such as chronic pain, depression, and sleep disorders, is one of the biggest problems in the United States. All these conditions are closely related, with about 35% of people with chronic pain also having comorbid depression and disordered sleep (Sun et al., 2025; Zhu et al., 2024). Given genetic data, the increased likelihood of developing major depressive disorder increases the chances of developing chronic localized pain and insomnia, with insomnia accounting for 12.8%-15.2% of the relationship between depression and pain (Zhu et al., 2024). Chronic pain decreases the quality of sleep and decreased sleep increases pain levels, as does chronic pain and insomnia. Depression also manifests as physical pain, which may increase insomnia exacerbating the pain-insomnia-depression triad (Satghare et al., 2022).


Presently, the clinical diagnostic paradigm is episodic and distinctly focused on one body system or type of problem, leaving clinicians dependent on patients' limited verbal self-reports (and often incomplete snapshots of their condition) to try and make informed clinical decisions. Therefore, a need

exists to implement continuous architecture-based monitoring as an alternative to episodic care across multiple classifications of care (Padliya, 2025a) as all other methods to clinically manage these individuals lead to a reactive approach to intervention with a significant delay in diagnosis.

We have developed the FUSE framework to address the issue stated above. This framework offers four guiding principles for a trustworthy multimodal diagnostic platform, which include 1. **Fidelity**: preserving the native temporal structure of each modality and using observed data. 2. **Uncertainty**: providing calibrated probabilities and conformal prediction sets. 3. **Synchronization**: syncing up cross-device and asynchronous data streams. 4. **Explain-Ability**: developing graph neighborhood narratives to link each of the aligned data points to their predictions. The FUSE framework has been operationalized using the Temporal Cross-Modal Alignment Graph (TCAG) methodology that fuses wearable with SMART phone and authorities with medical record data to establish a set of leading indicators for early detection of individuals who may experience a decline in condition within the pain-insomnia-depression triad.

Figure 1

The Pain-Insomnia-Depression Triad: A Bidirectional Feedback Loop



Literature Review

Applying artificial intelligence to digital health data has increased significantly. However, it is still learning to successfully navigate through three key challenges that the FUSE framework has addressed. These three challenges are: data paucity, data heterogeneity, and lack of clinical trust.

Data paucity presents one of the most significant challenges to the advancement of artificial intelligence in providing clinically meaningful information to healthcare providers. Even though digital sensors generate an unprecedented number of raw data points, quality clinical labels that can identify date/time indicators are rarely collected. Thus, self-supervised learning (SSL) has created great interest in its use; it enables AI technologies to develop clinically relevant representations from non-labeled data points. A review of 41 studies utilizing sensor-based technology identified that wearable platforms, such as wristbands containing gyroscopes, accelerometers, heart rate monitors, and voice analyzers, demonstrate high feasibility in monitoring depression. However, they provide no clinically accepted validation (Rahimi, 2026). By learning the structure of physiological time series, SSL models can use the finite clinical labels to be retrained for use as elements of scalable diagnostic systems.

Data heterogeneity limits all simplistic approaches to fusing disparate clinical data sets. Real-world patients exhibit multisensory modalities with asynchronous and incomplete data collections arriving at different sampling frequencies and temporal resolutions. Attempts to fuse patient data using naïve methods do not preserve the modality differences. A study using commercially available consumer-grade wrist trackers compared two mixed-input models with separate models that provided only time series activity data (without any corresponding information); the mixed-input models achieved a Hamming loss of 0.0946 for correctly classifying patients with comorbid depression and anxiety. The mixed models clearly outperformed the single modalities in classification accuracy (Lee et al., 2024). More sophisticated graph-based models will be able to explicitly model the temporal (lagged) relationships between modalities (i.e., there is a lagged effect of inadequate sleep on next-day mood) (Vasudevan et al., 2022). The use of such designs also allows for accommodating the cross-condition dependencies of the pain-insomnia-depression triad. Vasudevan et al. (2022) build a fixed graph over pre-aligned features. TCAG instead learns temporal alignment directly from raw, irregularly sampled streams before any graph is constructed, and couples this to calibrated uncertainty quantification at the output stage.

The lack of clinical trust represents the most significant barrier to the use of AI within healthcare. In a high-stakes situation, no clinician would put their faith in an AI-generated risk score without an explanation of that score (Murdoch, 2021). The issue of clinical trust

has received recognition from regulatory agencies. The FDA has called for open and transparent processes for determining the difference between AI-enabled medical devices, such as accountability infrastructures and bias management procedures (Shick et al., 2024). There has also been greater interest in explaining AI models to healthcare providers to allow clinicians to determine the underlying reasoning for why AI generated a risk score, as well as quantifying the model's uncertainty that created the risk score (Shahin et al., 2025). The FUSE framework attempts to resolve three key barriers (data paucity, data heterogeneity, and trust) in a manner that is clinically relevant to providers (using TCAG) for monitoring the triad of pain, insomnia, and depression.

Methods

Framework Objectives

The main aim is to clearly outline both FUSE and show how TCAG realized the four FUSE principles. The secondary aims are to: 1. Explain the preprocessing and data substrate for the pain-insomnia-depression triad, 2. Explain self-supervised pre-training strategy per modality, 3. Define TCAG compared to existing standard baseline models for fusion, and 4. Provide an evaluation plan to validate the results with pre-registered endpoints and fairness audits.

The FUSE Framework

FUSE defines four non-negotiable properties of multimodal diagnostic systems operating on clinical data. The principles are as follows.

Fidelity requires that each modality's native temporal structure be preserved through the full processing pipeline. Models must operate on the observed data graph. It should ensure data missingness is treated explicitly rather than concealed. Fidelity is instantiated in the TCAG data substrate and self-supervised encoding stages (Columns A and B of Figure 3).

Uncertainty requires that the system produce calibrated probability outputs accompanied by statistically rigorous prediction sets. Temperature scaling on a held-out validation split corrects for deep model overconfidence, and conformal prediction produces set-valued outputs that bound the error rate under weak distributional assumptions. Uncertainty is instantiated in the TCAG inference stage (Column D of Figure 3).


Synchronization requires that the system computationally account for clock offsets between devices, irregular sampling rates, and local temporal lags between modalities. The cross-modal alignment step of TCAG uses differentiable dynamic-time-warping-inspired alignments to handle these real-world imperfections (Column C of Figure 3).

Explainability means the system must show its reasoning in terms a clinician can understand. It should connect specific aligned events across modalities to the

final prediction, not just output a score. TCAG does this through graph-neighborhood explanations, which identify which modality windows contributed most to a given risk score. The result is a clinical narrative that a clinician can question and check, shown as (Column E in Figure 3).

Data from wearables, smartphones, and clinical surrogates are encoded and fused through the TCAG core. It produces calibrated risk scores, conformal prediction sets, and explainability narratives for clinician review. FUSE framework principles are mapped to corresponding pipeline stages.

Figure 2
End-to-End Platform Architecture Instantiating the FUSE Framework



Data Substrate and Preprocessing

The data substrate comprises three types of elements: physiologic time series; short-term acoustic samples; and tabular clinical surrogates. Physiology includes high-frequency streams created by devices worn by consumers, photoplethysmography (PPG) which measures pulse, and accelerometry (ACC) which measures physical activity as well as gait. It also includes SpO2 which measures oxygen levels in the blood and heart rate variability (HRV). Recent validations of consumer devices for capturing sleep data have indicated that they can create Hypnograms using these physiologic streams with an accuracy of greater than 90% for determining when a person is awake (Padliya, 2025b). Acoustic samples include recordings made by participants from coughing and completing speech tasks that were subsequently analyzed into Mel Spectrograms which capture information surrounding respiratory, emotional, and neurological states. Finally, clinical surrogates consist of realistic synthetic values for both laboratory tests and the unique matching ICD codes associated with those tests in chronological order; allowing for initial modeling prior to ever having access to real patient protected health information. All datasets, transformations, and models are versioned with immutable provenance to assure complete replication of results.

Self-Supervised Pretraining

A main element of this research's methodology was to employ self-supervised objectives to extract meaningful representations from a data stream without first being labeled with a clinical designation. To pretrain a model on physiological time series, masked-


step prediction and contrastive augmentation were applied, allowing the model to obtain invariance to motion artifacts, signal drift due to the device used to obtain the signal, and circadian cycles. The acoustic data was pretrained via patch-level instance discrimination on the spectrogram of the signal, providing detailed information on prosodic features, timbral characteristics of sound, and acoustic textures associated with respiratory and mental health disorders. The surrogate sequences used to develop the clinical models were then trained via time-aware transformers operating on the temporal trend of the surrogate sequences to extract latent states associated with each of the subjects in the data set. Each of the pretrained encoders produced a fixed-size embedding associated with each window of data, representing the data's native temporal structure and containing uncertainty.

Modeling Strategy: Baselines and TCAG Architecture

The comparative analysis evaluates the TCAG against two established baseline architectures. The Early Fusion baseline concatenates extracted features from all modalities into a single high-dimensional vector before classification. The Late Fusion baseline trains separate classifiers per modality and fuses predictions through a weighted average. While robust, neither baseline is optimized for asynchronous, often-missing real-world data streams. Early Fusion is brittle to missing data, and Late Fusion fails to capture cross-modal interactions at different timescales. The TCAG, introduced in Figure 3, is specifically designed to overcome these limitations.

Figure 3

The Temporal Cross-Modal Alignment Graph (TCAG) Architecture



Note. Five processing stages (A through E) implement the FUSE framework principles. The TCAG core (Column C, highlighted in orange) performs differentiable cross-modal alignment followed by heterogeneous graph construction and graph attention propagation. FUSE principle labels are shown beneath each column.

Methodological Core: Temporal Cross-Modal Alignment Graph

The TCAG fusion process is executed in two stages. The first stage is to find differentiable, dynamic-time-warping-based alignments between modalities, accounting for clock offsets between devices, non-uniform sampling rates, and local time lags, e.g., between the apneic event (SpO2 desaturation) and the ensuing arousal (HRV, ACC). Second, a heterogeneous graph of modality windows is built with the modality-specific pretrained encoders, where each node in the graph corresponds to a modality window. Edges in the graph correspond to two types of relationships between the nodes - temporal adjacency within the same data stream, and cross-modal affiliative relationships between streams. Edge weights between each of the two nodes are derived from the alignment score. Information is propagated through a graph attention means permitting reasoning about synchronous fast to synchronous slow events, as well as slower trends in context. Missingness of the data is also made explicit in observing subgraph (while ensuring real data gaps are satisfied with high fidelity) fulfilling the Fidelity Principle as instantiated by FUSE.

Inference, Calibration, and Explainability

The final inference stage maps the fused graph-based representation to a calibrated risk score. Temperature scaling on a held-out validation split corrects for the tendency of deep models to be overconfident, instantiating the Uncertainty principle of FUSE. Conformal prediction produces set-valued outputs with mathematically bounded error rates: the system can state it is "95% confident the patient's state is one of {mild, moderate} depression" rather than a point estimate. Explanations are generated at the level of graph neighborhoods and modality windows, linking specific aligned events to the model's final prediction, directly instantiating the Explainability principle.

Proposed Protocol for Prospective Clinical Validation

The following protocol is proposed for prospective validation. All primary and secondary endpoints must be pre-registered before conducting the empirical research studies. Saxena et al. (2025) found that approximately one-third (31%) of clinical GN (graph

neural network) published studies have not completed pre-registered analysis, demonstrating that pre-registered analysis is necessary for fixed data split use cases. Model discrimination will be reported via AUROC curves and will also be reported using a predetermined operating point to provide the additional metrics of sensitivity and specificity. For clinical decision-making, it is critical to select the operating point for the model with a sensitivity of 90% or higher to limit the number of false-negative (FN) predictions, even if this results in an increased number of false-positive (FP) predictions that are forwarded for review by clinicians. Model calibration will be assessed by using both reliability diagrams and Brier score decomposition. The net-benefit of deploying the model will be quantified using decision-curve analysis to compare against baseline and standard heuristic rules (DCA; Florkowski, 2008). Robustness testing will include stress-testing the model through synthetic dropout of sensors, misalignment of time, or shifting of domains. A fairness audit will report stratified performance based on demographics to ensure that the model does not perpetuate inequities.

Deployment and Interoperability

The purpose of this framework is for real-world use from the beginning. It ends with a microservice that provides timestamped information and low latency. For instance, a SMART on FHIR demonstrator would authenticate via OAuth2 and render the model's calibrated probabilities, conformal sets, and explanatory narratives inside a mock clinical chart (Mandel et al., 2016). A written threat model addressing re-identification risk, with mitigations including on-device preprocessing, federated training (Rehman et al., 2023), and differential privacy (Rocher et al., 2019), would be a prerequisite for real deployment.

Results

This section presents a structured architectural comparison of the Early Fusion baseline, the Late Fusion baseline, and the TCAG. It is evaluated across five dimensions directly relevant to clinical deployment of multimodal diagnostic AI (see Table 1). The comparison draws on properties established in the Methods section and on published evidence.

Table 1
Architectural Comparison Over Deployment Dimensions

Missing data handling	imputes full feature vector when modality is absent	Partial; excludes absent modality	Explicit; uses only observed subgraph
Cross-modal interaction	Limited to feature concatenation only	Not captured; modalities processed in isolation	Captured via heterogeneous graph edges
Temporal alignment	None; assumes synchronous streams at identical rates	Limited; aligns each modality internally	Differentiable DTW-inspired alignment
Uncertainty quantification	Not integrated	not calibrated	Calibrated via temperature scaling
Explainability	Post-hoc attribution over concatenated feature vector	Per-modality attribution possible; no cross-modal narrative	Graph-neighborhood explanations linking aligned events to predictions

Table 1 demonstrates that TCAG addresses all five deployment-critical dimensions, whereas Early Fusion and Late Fusion each fail on multiple counts. The most clinically consequential gap is missing data handling. Sensor streams break down in the real world. A device goes unworn. A battery dies. A patient skips a reading. Early Fusion collapses under these gaps. Late Fusion survives them but loses the cross-modal signal. TCAG takes a different approach. Its subgraph method never fabricates a value to fill a gap. It works only with what was observed. This is the Fidelity principle of FUSE in practice. The integrated uncertainty quantification via conformal prediction addresses the regulatory demand for transparent AI-enabled devices (Shick et al., 2024), providing set-valued outputs with mathematically guaranteed coverage rather than uncalibrated point estimates.

Empirical Results: NHANES Cross-Sectional Analysis

In addition to examining the architectural comparison above, a calibrated gradient-boosting

classifier was employed to model data from the NHANES 2017-2018 data set that included demographic information as well as data related to PHQ-9 depression screens, sleep, physical activity, cardiac/metabolic history, and a chronic pain functional limitation proxy using information from the Physical Functioning questionnaire (n = 3,053 U.S. adults with complete datasets). The classifier yielded an AUROC of 0.821 for the primary outcome of interest (PHQ-9 score ≥ 10 ; 8.3% tested prevalence), and exhibited reasonable calibration based on Platt scaling (Brier = 0.066) and split conformal prediction achieving 89.7% empirical coverage. The chronic pain functional limitation proxy represented the strongest predictor accounting for 45.5% of the feature importance, which was over two times greater than the second strongest predictor (age, 17.0%) consistent with the triadic relationship that served as motivation for the article. Analysis was conducted cross-sectionally with reliance on self-reported data; thus, while the analysis indicates that a triadic relationship exists, it does not provide clinical validation of TCAG.

Figure 4
Receiver Operating Characteristic Curve for NHANES Depression Classification




Figure 5
Calibration (Reliability) Diagram for the NHANES Classifier





Figure 6
Feature Importance for NHANES Depression Classification



Empirical Results: TCAG Synthetic Architecture Evaluation

A second analysis used synthetic data with known ground truth to isolate TCAG's asynchronous-fusion behavior, which the NHANES analysis cannot exercise. A binary deterioration-event label depended on a coordinated, lagged signature spanning three simulated modalities with independent clock offsets. No single modality alone carried a detectable signal. An initial single-seed comparison produced unstable results and was discarded; all results below average over multiple independent random seeds. In the baseline

comparison, TCAG achieved the highest mean discrimination (AUROC = 0.839, SD = 0.015) versus Early Fusion (0.779, SD = 0.043) and Late Fusion (0.741, SD = 0.048). Under increasing sensor dropout (0-60%), TCAG led at every level and degraded most gracefully (0.713 vs. 0.631 and 0.584 at 60% dropout). Under increasing temporal misalignment (0-10 minutes), TCAG again led at every level and was markedly more stable (0.878 to 0.848) than Late Fusion (0.846 to 0.712 to 0.723). These are architecture-characterization results on synthetic data, not a prediction of performance on real clinical data.

Figure 7

AUROC Under Increasing Sensor Dropout (Mean $\pm$ SD Across Seeds)




Figure 8

AUROC Under Increasing Temporal Misalignment (Mean $\pm$ SD Across Seeds)





Figure 9
Qualitative Trace of Learned Cross-Modal Attention



Interpretation and Clinical Significance

The implementation of the FUSE framework and TCAG architecture is likely to significantly change how clinicians manage the pain-insomnia-depression triad and how clinicians use health AI to deliver multimodal healthcare. Once evaluated through further empirical studies, the FUSE Framework and TCAG Architecture will give clinicians real-time insight into patient data regarding the delivery of care in relation to the pain-insomnia-depression triad, thereby enabling clinicians to shift their approach to care from a reactive and fragmented approach to a proactive and integrated approach to care. In relation to the use of TCAG Architecture, clinicians could receive alerts when their patients' conditions deteriorate, thus enabling them to take appropriate steps to avert costly hospitalizations by intervening earlier. Moreover, the use of TCAG Architecture as a framework for working with other complex multimodal diagnostic issues in medicine will provide a roadmap for healthcare providers dealing with such issues.

Deployment Considerations

The implementation of this form of translational research faces multiple challenges before it can be translated into health care practice. Currently, healthcare patient data are often siloed and incompatible across the various electronic health record (EHR) systems available today; the SMART on FHIR platform is one approach to help overcome this challenge (Mandel et al., 2016). Continuous sensor data will be required to train deep-learning models; to obtain sufficient computational resources to train such models requires large-scale distribution of hardware and software, which can be achieved using federated learning and self-supervised pretraining techniques (Rehman et al., 2023). There are ethical issues associated with using voice and gait data to identify individuals. Both voice and gait could be used to uniquely identify an individual in a manner as conclusively as a fingerprint (Rocher et al., 2019). Thus, creating a successful operational model of this type will necessitate establishing a rigorous data governance framework with clearly defined patient consent processes and clearly articulated agreements governing how the information collected will be utilized (Murdoch, 2021).

Strengths and Limitations

This work combines an analysis of actual data from a large, representative sample (NHANES) with a simulated synthetic evaluation of the proposed architecture with known ground truth. The NHANES analysis was retrospective (cross-sectional) and dependent upon self-reports (sleep and functional limitation) instead of objective sensors. For TCAG, the evaluation was conducted exclusively using synthetic data derived from a highly simplified generative structure.

Future Research Directions

Priority directions include re-evaluating depression prediction using objective wearable sensor data rather than self-report and evaluating TCAG on real asynchronous multimodal clinical data. The Medical Information Mart for Intensive Care (MIMIC-IV) database is a natural candidate given its continuous, irregularly sampled structure, and such an analysis would be direct empirical follow-up to this paper. This work has focused on demonstrating feasibility for the pain-insomnia-depression triad. The architecture is also expected to extend to neurological disorders using passive gait and voice biomarkers. It can also be extended to other multi-condition comorbidity clusters where continuous monitoring is clinically superior to episodic monitoring methods.

Conclusion

This paper presented the FUSE framework (Fidelity, Uncertainty, Synchronization, and Explainability) for building reliable multimodal diagnostic AI in digital health. The Temporal Cross-Modal Alignment Graph (TCAG) is the architecture that puts FUSE into practice, combining wearable, mobile, and electronic health record data to study the relationship between chronic pain, depression, and sleep disruption. TCAG represents these interactions as a graph, handling missing data through subgraph analysis, uncertainty through conformal prediction, cross-modal timing through differentiable alignment, and explanation through neighborhood-based graph reasoning. A synthetic evaluation showed that TCAG outperformed Early and Late Fusion baselines (AUROC = 0.839 vs. 0.779 and 0.741). A cross-sectional analysis of NHANES survey data (AUROC = 0.821) showed that chronic-pain functional limitation is a strong, independent correlate and supports the clinical premise behind this work. The goal of this research is to give clinicians a decision support tool for complex chronic conditions to aid in their clinical judgment.

Acknowledgement

The author declares no external funding was received for this work and no conflicts of interest to report.

References

- Florkowski, C. M. (2008). Sensitivity, specificity, positive and negative predictive values for diagnostic tests. *The Clinical Biochemist Reviews*, 29(Suppl 1), S81–S86.
<https://pmc.ncbi.nlm.nih.gov/articles/PMC2556590/>
- Lee, T., Kim, G. H., & Choi, M. (2024). Geriatric depression and anxiety screening via deep learning using activity tracking and sleep data. *International Journal of Geriatric Psychiatry*, 39(2).
<https://doi.org/10.1002/gps.6071>
- Mandel, J. C., Kreda, D. A., Mandl, K. D., Kohane, I. S., & Ramoni, R. B. (2016). SMART on FHIR: A standards-based, interoperable apps platform for electronic health records. *Journal of the American Medical Informatics Association*, 23(5), 899–908.
<https://doi.org/10.1093/jamia/ocv189>
- Murdoch, B. (2021). Privacy and artificial intelligence: Challenges for protecting health information in a new era. *BMC Medical Ethics*, 22(1), 115.
<https://doi.org/10.1186/s12910-021-00687-3>
- Padliya, T. (2025a). Zap the pain, skip the pills with hypersound: Noninvasive chronic pain treatment with a new spin on radio frequency microwaves. *IEEE Pulse*, 16(2), 22–25.
<https://doi.org/10.1109/MPULS.2025.3554963>
- Padliya, T. (2025b, July/August). The new night watch: Advanced sleep sensors uncover what keeps us awake. *IEEE Pulse*, 16(4), 34–40.
<https://doi.org/10.1109/MPULS.2025.3598836>
- Powell, D. (2024). Walk, talk, think, see and feel: Harnessing the power of digital biomarkers in healthcare. *npj Digital Medicine*, 7(1), 116.
<https://doi.org/10.1038/s41746-024-01023-w>
- Rahimi, M., Abrishamifar, K., Hazhir, S., Valizadeh, H., Nourani, A., & Rahimi, B. (2026). Sensor-equipped digital technologies for monitoring and detecting depressive disorders: A systematic review. *Health Science Reports*, 9(3).
<https://doi.org/10.1002/hsr2.71743>
- Rehman, M. H. U., Dirir, A. M., Salah, K., Svetinovic, D., & Bentahar, J. (2023). Federated learning in medical imaging: A systematic review. *The British Journal of Radiology*, 96(1152), 20220890.
<https://doi.org/10.1259/bjr.20220890>
- Rocher, L., Hendrickx, J. M., & de Montjoye, Y. A. (2019). Estimating the success of re-identifications in incomplete datasets using generative models. *Nature Communications*, 10(1), 3069.
<https://doi.org/10.1038/s41467-019-10933-3>
- Satghare, P., Abdin, E. B., Hombali, A., Lin Teh, W., Samari, E., Chua, B. Y., Verma, S., Mok, Y. M., Chong, S. A., Subramaniam, M., & Ushida, T. (2022). Chronic pain among tertiary care psychiatric out-patients in Singapore: Prevalence and associations with psychiatric disorders. *Pain Research and Management*, 2022, Article 1825132.
<https://doi.org/10.1155/2022/1825132>
- Saxena, A., Nixon, B., Boyd, A., Evans, J., & Faraone, S. V. (2025). A systematic review of the application of graph neural networks to extract candidate genes and biological associations. *American Journal of Medical Genetics Part B: Neuropsychiatric Genetics*, 198(6), 3–18.
<https://doi.org/10.1002/ajmg.b.33031>
- Shahin, M. H., Desai, P., Terranova, N., Guan, Y., Helikar, T., Lobentanzer, S., Liu, Q., Lu, J., Madhavan, S., Mo, G., Musuamba, F. T., Podichetty, J. T., Shen, J., Xie, L., Wiens, M., & Musante, C. J. (2025). AI-driven applications in clinical pharmacology and translational science: Insights from the ASCPT 2024 AI Preconference. *Clinical and Translational Science*, 18(4).
<https://doi.org/10.1111/cts.70203>
- Shick, A. A., Webber, C. M., Kiarashi, N., Weinberg, J. P., Deoras, A., Petrick, N., Saha, A., & Diamond, M. C. (2024). Transparency of artificial intelligence/machine learning-enabled medical devices. *npj Digital Medicine*, 7, 23.
<https://doi.org/10.1038/s41746-023-00992-8>
- Sun, H., Chen, P., Bai, W., Feng, Y., Sha, S., Su, Z., Cheung, T., Ng, C. H., Zhang, Q., Xiang, Y., & Zhong, B. (2025). Depressive and insomnia symptoms among older adults with different chronic pain trajectories: A network analysis based on observation over an eight-year period. *Depression and Anxiety*, 2025, Article 8065167.
<https://doi.org/10.1155/da/8065167>
- Vasudevan, S., Srinivasan, R., & Jagannathan, S. (2022). Digital health and biomarkers: The convergence. *npj Digital Medicine*, 5(1), 36.
<https://doi.org/10.1038/s41746-022-00583-z>
- Zhu, Y., Bi, Y., & Zhu, T. (2024). Mendelian randomization highlights sleep disturbances mediated the effect of depression on chronic pain. *Brain and Behavior*, 14(7).
<https://doi.org/10.1002/brb3.3596>