

VitalDiagnosis: AI-Driven Ecosystem for 24/7 Vital Monitoring and Chronic Disease Management

Zhikai Xue¹, Tianqianjin Lin², Pengwei Yan², Ruichun Wang¹
Yuxin Liu³, Zhuoren Jiang^{2*}, Xiaozhong Liu^{1*}

¹Worcester Polytechnic Institute

²Zhejiang University

³University of Delaware

{zxue1, rwang12, xliu14}@wpi.edu, {lintqj, yanpw, jiangzhuoren}@zju.edu.cn, cindyliau@udel.edu

Abstract

Chronic diseases have become the leading cause of death worldwide, a challenge intensified by strained medical resources and an aging population. Individually, patients often struggle to interpret early signs of deterioration or maintain adherence to care plans. In this paper, we introduce VitalDiagnosis, an LLM-driven ecosystem designed to shift chronic disease management from passive monitoring to proactive, interactive engagement. By integrating continuous data from wearable devices with the reasoning capabilities of LLMs, the system addresses both acute health anomalies and routine adherence. It analyzes triggers through context-aware inquiries, produces provisional insights within a collaborative patient-clinician workflow, and offers personalized guidance. This approach aims to promote a more proactive and cooperative care paradigm, with the potential to enhance patient self-management and reduce avoidable clinical workload.

Demo Repository — <https://tinyurl.com/5n83hcrz>

Introduction

Chronic diseases, including cardiovascular conditions, diabetes, and stroke, have become the leading causes of death worldwide (Stuckler 2008; Strong et al. 2005). Unlike common illnesses, managing chronic diseases is a long-term, complex process that requires sustained medical engagement and considerable resource allocation, contributing to growing healthcare demands (Wagner 1998; Wagner, Austin, and Von Korff 1996; Tan et al. 2024). This pressure, intensified by an aging population, has strained medical systems and led to delays in timely care delivery (Smith 2017), limiting opportunities for early intervention that could mitigate severe complications (Adigun et al. 2019). At the individual level, patients often lack the clinical literacy needed to self-manage their conditions or recognize early signs of deterioration (Lu et al. 2025). As a result, healthcare systems remain largely reactive, struggling to keep pace with the escalating burden of chronic illness (Choi et al. 2024; GBD 2019 ACC Collaborators 2025).

To address these challenges, we introduce **VitalDiagnosis**, an ecosystem of specialized LLMs designed to shift

chronic care from passive monitoring to proactive engagement. The entire workflow is powered by Unified Memory Core, an architecture in which a central Memory MiniLLM maintains deep, continuous context by integrating medical knowledge bases and patient-related assets with an adaptable parametric memory composed of shared and personalized LoRAs. The process begins with a lightweight multimodal Monitoring MiniLLM that interprets wearable data to produce clinically relevant narratives and identify potential trigger events. The system then activates its dual-track framework through a Domain LLM specialized in clinical inquiry: for potential outliers, it initiates a dialogue to contextualize the event; for routine management, it engages patients to monitor adherence. The resulting insights are synthesized into provisional clinical assessments and actionable recommendations, which are circulated within a collaborative patient-clinician workflow for timely review. This proactive, data-driven approach aims to promote more anticipatory chronic care, with the potential to enhance self-management and reduce avoidable clinical workload.

Related Work

While LLMs show significant potential in healthcare (Goyal et al. 2024; Li et al. 2025; Zhou, Song, and Shen 2025b), their integration with wearable devices remains nascent (Fang et al. 2024; Ferrara 2024; Zhang et al. 2025). Prior work in wearable-based chronic care has largely been confined to passive, threshold-based systems that trigger simple alerts for outliers (de Souza Ferreira et al. 2024; Xie et al. 2021; Subramanian et al. 2020). These approaches fundamentally lack the capacity for nuanced, interactive investigation of anomalies and fail to provide continuous, adaptive support for personalized routine adherence (Zhou, Song, and Shen 2025a). In contrast, our VitalDiagnosis ecosystem addresses these gaps by introducing a dual-track framework that leverages the synergy of LLMs and wearables to support both real-time, interactive triage and proactive routine care within a collaborative patient-clinician loop.

VitalDiagnosis Ecosystem

VitalDiagnosis integrates LLMs with wearables into a dual-track framework for chronic disease management. By supporting both interactive triage of anomalies and proac-

*Corresponding authors.

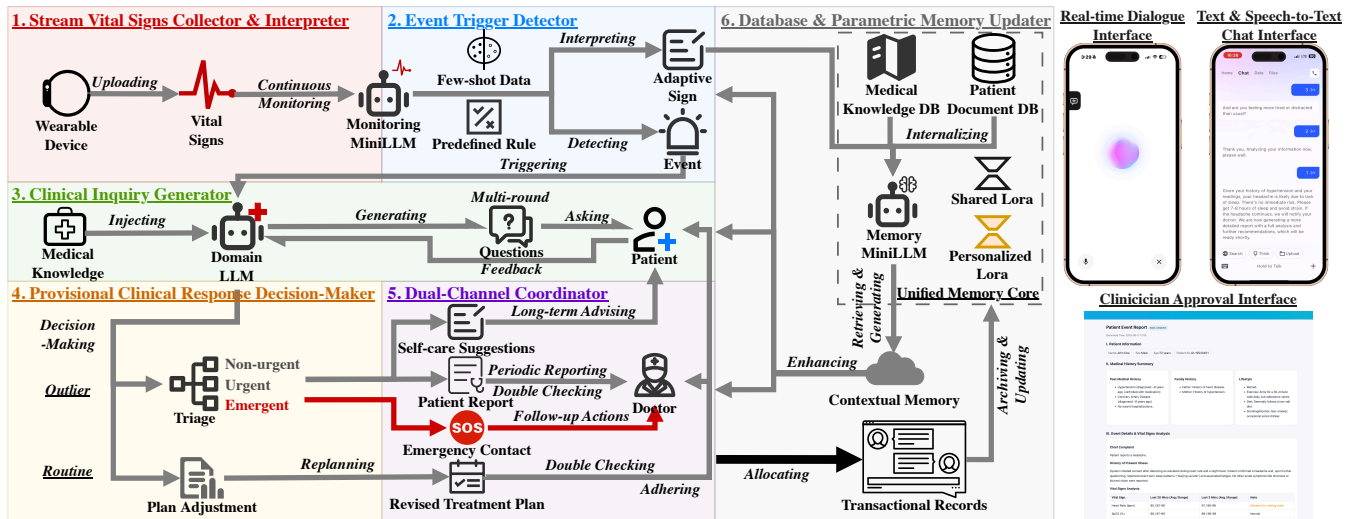


Figure 1: The framework and interfaces of the VitalDiagnosis ecosystem.

tive adherence monitoring, it establishes a collaborative patient–clinician loop that promotes more anticipatory care. Powering this workflow is the Unified Memory Core. Mediated by a central Memory MiniLLM, it combines persistent medical-knowledge databases and patient assets with a parametric memory of shared and personalized LoRAs (Zhang, Kim, and Liu 2024), producing episodic or structured summaries that condition all downstream processes.¹

Stream Vital Signs Collector & Interpreter The workflow begins by continuously collecting vital sign streams from wearable devices. A lightweight multimodal Monitoring MiniLLM then interprets these raw, variable-length signal segments into concise, clinician-readable narratives, supplying the initial context for all subsequent tasks.

Event Trigger Detector Next, the system identifies clinically relevant triggers by inferring the patient’s state using track-specific logic. It applies rule-based thresholds and model-based inference for anomalies, while also scheduling periodic checks for routine care. The system then routes each event to either *outlier-detection* or *routine-adherence* logic, producing a risk-graded trigger to guide the next stage.

Clinical Inquiry Generator Upon event triggering, a scene-specialized Domain LLM, adapted with LoRA on clinician-annotated simulated and rewritten cases, initiates a brief Q&A session with the patient, conditioned on contextual memory. For outlier events, it explores symptoms and contributing factors; for routine management, it assesses adherence and notes potential barriers. The inquiry remains minimal, concluding once adequate information is obtained.

Provisional Clinical Response Decision-Maker Using interpreted events, inquiry results, and contextual memory, the system produces a *provisional* clinical response. For outliers, it provides severity-aware triage with potential factors

¹We utilize LLMs of varying sizes: Memory MiniLLM (4B), Monitoring MiniLLM (1.7B), and Domain LLM (14B).

and next steps; for routine care, it summarizes adherence and suggests refinements. All recommendations follow a tiered approval process, with high-risk items requiring explicit review and low-risk items eligible for deferred confirmation.

Dual-Channel Coordinator The system then coordinates all communication, translating provisional judgments into audience-specific reports bounded by visibility constraints. It delivers triage-aware guidance to patients while presenting providers with a concise event summary and any flagged high-risk proposals for review. For routine management, it offers patients adherence suggestions and supplies providers with periodic digests for confirmation.

Database & Parametric Memory Updater To complete the loop, the system updates its memory with all interactions and outcomes. The memory database maintains a short-term rolling snapshot of recent activity, while confirmed facts are promoted to long-term storage with full traceability. The parametric memory (LoRA modules) is updated more selectively on stable patterns emerging from the long-term data, ensuring the system is continuously refined through accumulated interactions and expert oversight.

Conclusion

In conclusion, VitalDiagnosis presents an ecosystem that integrates LLMs with wearables to support a dual-track framework for chronic disease management. By enabling both interactive triage of health anomalies and proactive monitoring of routine adherence, it fosters a collaborative patient–clinician loop that promotes more anticipatory care. This scalable, proactive paradigm has the potential to enhance patient self-management, ease provider workload, and broaden access to continuous, personalized support. To further assess its real-world performance and contribute to the community, VitalDiagnosis is currently undergoing pilot studies with medical institutions, and a curated, clinician-annotated dataset will be made publicly available.

References

- Adigun, A. C.; Maguire, K.; Jiang, Y.; Qu, H.; and Austin, S. 2019. Urgent care center and emergency department utilization for non-emergent health conditions: analysis of managed care beneficiaries. *Population Health Management*, 22(5): 433–439.
- Choi, M.; Sempungu, J. K.; Lee, E. H.; and Lee, Y. H. 2024. Living longer but in poor health: healthcare system responses to ageing populations in industrialised countries based on the Findings from the Global Burden of Disease Study 2019. *BMC Public Health*, 24(1): 576.
- de Souza Ferreira, E.; de Oliveira, A. H. M.; Dias, M. A.; da Costa, G. D.; Januário, J. P. T.; Botelho, G. M.; and Cotta, R. M. M. 2024. Mobile solution and chronic diseases: development and implementation of a mobile application and digital platform for collecting, analyzing data, monitoring and managing health care. *BMC Health Services Research*, 24(1): 1009.
- Fang, C. M.; Danry, V.; Whitmore, N.; Bao, A.; Hutchison, A.; Pierce, C.; and Maes, P. 2024. Physiollm: Supporting personalized health insights with wearables and large language models. In *2024 IEEE EMBS International Conference on Biomedical and Health Informatics (BHI)*, 1–8. IEEE.
- Ferrara, E. 2024. Large language models for wearable sensor-based human activity recognition, health monitoring, and behavioral modeling: a survey of early trends, datasets, and challenges. *Sensors*, 24(15): 5045.
- GBD 2019 ACC Collaborators. 2025. Characterising acute and chronic care needs: insights from the Global Burden of Disease Study 2019. *Nature communications*, 16(1): 4235.
- Goyal, S.; Rastogi, E.; Rajagopal, S. P.; Yuan, D.; Zhao, F.; Chintagunta, J.; Naik, G.; and Ward, J. 2024. Healai: A healthcare llm for effective medical documentation. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, 1167–1168.
- Li, C.; Zhao, Y.; Bai, Y.; Zhao, B.; Tola, Y. O.; Chan, C. W.; Zhang, M.; and Fu, X. 2025. Unveiling the potential of large language models in transforming chronic disease management: mixed methods systematic review. *Journal of Medical Internet Research*, 27: e70535.
- Lu, T.; Lin, Q.; Yu, B.; and Hu, J. 2025. A systematic review of strategies in digital technologies for motivating adherence to chronic illness self-care. *npj Health Systems*, 2(1): 13.
- Smith, S. M. 2017. Disparities in emergency and urgent care services in rural and urban communities. *Papers & Publications: Interdisciplinary Journal of Undergraduate Research*, 6(1): 9.
- Strong, K.; Mathers, C.; Leeder, S.; and Beaglehole, R. 2005. Preventing chronic diseases: how many lives can we save? *The Lancet*, 366(9496): 1578–1582.
- Stuckler, D. 2008. Population causes and consequences of leading chronic diseases: a comparative analysis of prevailing explanations. *The Milbank Quarterly*, 86(2): 273–326.
- Subramanian, M.; Wojtusciszyn, A.; Favre, L.; Boughorbel, S.; Shan, J.; Letaief, K. B.; Pitteloud, N.; and Chouchane, L. 2020. Precision medicine in the era of artificial intelligence: implications in chronic disease management. *Journal of translational medicine*, 18: 1–12.
- Tan, S. Y.; Sumner, J.; Wang, Y.; and Wenjun Yip, A. 2024. A systematic review of the impacts of remote patient monitoring (RPM) interventions on safety, adherence, quality-of-life and cost-related outcomes. *NPJ Digital Medicine*, 7(1): 192.
- Wagner, E. H. 1998. Chronic disease management: what will it take to improve care for chronic illness? *Effective clinical practice*, 1(1).
- Wagner, E. H.; Austin, B. T.; and Von Korff, M. 1996. Organizing care for patients with chronic illness. *The Milbank Quarterly*, 511–544.
- Xie, Y.; Lu, L.; Gao, F.; He, S.-j.; Zhao, H.-j.; Fang, Y.; Yang, J.-m.; An, Y.; Ye, Z.-w.; and Dong, Z. 2021. Integration of artificial intelligence, blockchain, and wearable technology for chronic disease management: a new paradigm in smart healthcare. *Current Medical Science*, 41(6): 1123–1133.
- Zhang, K.; Kim, Y.; and Liu, X. 2024. Personalized llm response generation with parameterized memory injection. *arXiv preprint arXiv:2404.03565*.
- Zhang, Y.; Ayush, K.; Qiao, S.; Heydari, A. A.; Narayanswamy, G.; Xu, M. A.; Metwally, A. A.; Xu, S.; Garrison, J.; Xu, X.; et al. 2025. SensorLM: Learning the Language of Wearable Sensors. *arXiv preprint arXiv:2506.09108*.
- Zhou, Y.; Song, L.; and Shen, J. 2025a. Improving Medical Large Vision-Language Models with Abnormal-Aware Feedback. In Che, W.; Nabende, J.; Shutova, E.; and Pilehvar, M. T., eds., *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 12994–13011. Vienna, Austria: Association for Computational Linguistics. ISBN 979-8-89176-251-0.
- Zhou, Y.; Song, L.; and Shen, J. 2025b. MAM: Modular Multi-Agent Framework for Multi-Modal Medical Diagnosis via Role-Specialized Collaboration. In Che, W.; Nabende, J.; Shutova, E.; and Pilehvar, M. T., eds., *Findings of the Association for Computational Linguistics: ACL 2025*, 25319–25333. Vienna, Austria: Association for Computational Linguistics. ISBN 979-8-89176-256-5.