

Tiered Desensitization and Lightweight Fusion of Multi-Source Heterogeneous Data for Privacy-Preserving Personalized Health Management in Community Elderly Care: A Case Study of the "Hui-YangLao" Platform in Chengdu, China

Yuting Liu¹, Huan Wen¹, Huan He²

¹ School of Electronic Commerce, Sichuan Mayflower Vocational and Technical College, Chengdu, China

² School of Culture and Communication, Sichuan Culture and Communication College, Chengdu, China

* Corresponding author: [3091587071@qq.com]

Funding: This work was supported by the 2026 Special-Feature Research Grant of the Hebei Research Association for Financial Services to Science and Technology Innovation (Grant No. JRFW2026017).

Abstract:Background: Community-based smart elderly-care platforms continuously aggregate wearable, primary-care, and community-service data about older adults. These data are highly sensitive, yet their granularity is exactly what personalized chronic-disease management and home-safety warning require. Resource-constrained county and community platforms, in particular, struggle to reconcile privacy compliance, computational cost, and service accuracy. Methods: We report a 10-month prospective deployment study on the "Hui-YangLao" platform in Pidu District, Chengdu, China, covering 1,267 older adults across six urban communities and three townships. We designed and operationalized a tiered desensitization scheme (high / medium / low privacy risk) coupled with a three-stage lightweight fusion pipeline (data-level cleaning, feature-level weighted fusion, decision-level cross fusion) and a four-tier role-based access control. We evaluated privacy protection, data utility, and computational overhead before and after the intervention, using platform-monitored re-identification risk, risk-identification accuracy, and server CPU/memory consumption. Results: After desensitization, the composite identity re-identification risk score dropped from 9.2/10 to 1.1/10 (88.0% reduction); geolocation re-identification probability fell from 100% to 0.7%; the coefficient of determination for chronic-disease trend fitting decreased only marginally ($R^2 = 0.961$ to 0.948). The chronic-risk identification accuracy dropped from 92.36% to 90.82% (1.54 percentage points) while the share of sensitive fields in cloud-stored records fell from 41.7% to 0%. Per-batch computational overhead decreased by 61.7%, allowing the pipeline to run on a low-cost community cloud server. Across the 10-month pilot, 17 unauthorized-access alerts were logged and closed within 24 hours, and no confirmed exfiltration of identifiable elderly health records occurred. Conclusions: For resource-constrained community elderly-care platforms, coupling tiered desensitization with a lightweight three-stage fusion pipeline can substantially reduce privacy exposure without materially degrading the accuracy of personalized health warnings. The scheme offers a low-cost, reproducible reference implementation for other mid-sized Chinese community platforms.

Keywords: smart elderly care; multi-source heterogeneous data; tiered desensitization; lightweight data fusion; privacy-preserving health management; community-based care; role-based access control

1. Introduction

China is aging faster than most countries in the world: by the end of 2023, adults aged 60 and older

numbered 297 million, accounting for 21.1% of the national population [5,6]. Policy has responded by promoting "aging in place" supported by information and communication technologies, wearable sensors, and community service platforms [7,8]. Smart bands, millimeter-wave sleep radars, gas and door/window sensors, primary-care electronic medical records (EMRs), and community care work orders now continuously produce large volumes of multi-source heterogeneous health data about older adults [9,10,34].

Two regulatory trends make this data landscape legally and ethically salient. China's Personal Information Protection Law (2021) and Data Security Law (2021) classify health diagnosis, treatment records, and precise home location as sensitive personal information, imposing strict rules on collection, processing, and cross-organization sharing [33]. At the same time, personalized chronic-disease assessment, fall detection, and home-safety warning depend precisely on continuous, fine-grained features. Over-aggressive anonymization destroys the variables these models need; under-protected sharing exposes older adults to re-identification, discrimination, and legal liability for platform operators [14,13].

Most existing privacy-preserving health-data research has been developed for well-resourced hospitals or large health-tech platforms, where a dedicated data middle-end, professional DevOps team, and server clusters are available [27,15]. County, sub-district, and township platforms in China, however, typically operate on a modest cloud budget, rely on part-time maintenance, and face unclear responsibility boundaries among civil affairs, health commissions, community committees, and vendor firms [32,30]. Directly porting hospital-grade fusion and encryption stacks to these settings is usually unaffordable and operationally fragile.

Two gaps remain in the literature. First, computer-science studies of privacy-preserving health-data publishing largely advance algorithmic mechanisms (k-anonymity [1], l-diversity [2], differential privacy [3,4], federated learning) but rarely report how these mechanisms are actually configured and tuned in a low-resource community deployment. Second, social-science and policy studies of Chinese smart elderly care describe service models, value co-creation, and governance at a macro level [31,36], but stop short of operational, field-tested data-governance procedures. The result is a gap between theoretically strong privacy methods and the messy engineering reality of small community platforms.

This paper addresses that gap through a 10-month prospective deployment study on the "Hui-YangLao" platform in Pidu District, Chengdu. We make three contributions. (1) We operationalize a tiered desensitization scheme that maps heterogeneous elderly-care data into high, medium, and low privacy-risk tiers and applies proportionate protection to each. (2) We couple this scheme with a lightweight three-stage fusion pipeline (data-level cleaning, feature-level weighted fusion, decision-level cross fusion) designed to run on low-cost community hardware. (3) We evaluate the deployment on three dimensions simultaneously: privacy protection, data utility for health warning, and computational cost, and we document the accompanying four-tier role-based access control and closed-loop alerting governance.

2. Related Work

2.1 Privacy-preserving health-data publishing

De-identification of health data has moved through three generations of methods. Sweeney's k-anonymity [1] requires each quasi-identifier group to contain at least k indistinguishable records; Machanavajjhala et al. [2] extended this with l-diversity to prevent attribute disclosure within a group. Dwork [3] introduced differential privacy (DP), which adds calibrated noise to query outputs and provides a formal guarantee independent of the attacker's background knowledge [4,20]. Subsequent work has shown that syntactic anonymity and DP each have trade-offs: DP offers stronger guarantees but at the cost of utility, while k-anonymity preserves utility but remains vulnerable to linkage attacks without proper constraints [19,18]. Recent health-data studies therefore combine mechanisms — for example, microaggregation-based k-anonymity to reduce DP noise [17], fine-grained randomized protection for sensitive elements while preserving utility through graphical models [15], or multi-attribute hierarchical DP that allocates privacy budgets according to sensitivity [16].

For ambient assisted living (AAL) and wearable health monitoring specifically, Psychoula et al. [11] used an LSTM encoder-decoder to generate multiple data views according to access level; Jourdan et al. [12] shifted protection to the wearable smartphone and showed that activity recognition could be preserved

while re-identification was substantially reduced. Chowdhury et al. [13] proposed a patient-controlled framework for real-time wearable data sharing. These studies establish the technical feasibility of tiered protection, but most are evaluated on benchmark datasets or simulated AAL environments, not on a year-long community deployment serving frail older adults.

2.2 Multi-sensor data fusion for elderly monitoring

Multi-sensor fusion is conventionally organized into three levels: data-level, feature-level, and decision-level [21,24]. Data-level fusion preserves raw signal resolution but demands bandwidth and compute; feature-level fusion compresses signals into representative features and is robust to heterogeneous sampling rates; decision-level fusion combines independent classifier outputs and is the most bandwidth-efficient [21]. Gravina et al. [21] reviewed multi-sensor fusion in body sensor networks and noted that the optimal level depends on available hardware, latency, and reliability requirements. For fall detection and activity recognition, recent work reports high accuracy on public benchmarks such as UTD-MHAD, but typically relies on deep models trained on high-end servers [26,22,37]. Nweke et al. [24], Muhammad et al. [23], and John et al. [25] reviewed the field and noted that fall detection in particular has been intensively studied [29]; they collectively conclude that practical deployment on low-cost edge/cloud hardware remains an open problem, particularly for rural and community settings.

2.3 Smart elderly-care platforms in China

Smart home-based elderly care in China has grown rapidly since 2008 and is largely policy-driven rather than demand- or technology-driven [6]. Recent platforms integrate wearable devices, community health records, and service work orders [8,9]. Wei et al. [9] analyzed a western-China community platform and found that the oldest-old and women living alone were the most dependent on smart-care services. Sun et al. [10], in an ethnographic study of community care in China, argued that technology should play a supportive rather than substitutive role and that privacy and trust are central to uptake. Hanpo et al. [8] and Wang and Xu [7] documented the architecture of "Internet+" community platforms but did not examine the data-governance layer. Public authorized data operation has been proposed as a structural solution to home-based smart elderly-care dilemmas in China [35], and integrated medical-nursing-rehabilitation care continues to drive demand for interoperable data infrastructure [28]. To our knowledge, no published study reports a field-tested, computationally lightweight desensitization-and-fusion pipeline for these under-resourced Chinese community platforms, with simultaneous evaluation of privacy, utility, and cost.

3. Materials and Methods

3.1 Study setting and the Hui-YangLao platform

This study was conducted on the "Hui-YangLao" (literally "Beneficial Elderly Care") platform, a community-based smart elderly-care system operated in Pidu District, Chengdu, Sichuan, China. The platform integrates (i) wearable and ambient sensing terminals, (ii) two community health-center EMR databases, and (iii) nine community care-station work-order systems. The pilot ran for 10 months, covering six urban communities and three townships. Ethical approval was obtained from the [name of ethics committee, approval number to be inserted] prior to deployment; written informed consent was obtained from each participant (or their legal representative for older adults with cognitive impairment), and all data were processed under the platform's data-sharing agreement with the participating community health centers.

3.2 Data sources and sample

A total of 1,267 community-dwelling older adults were enrolled: 582 aged 60–69, 465 aged 70–79, and 220 aged 80 and above (including 220 classified as moderately or severely disabled). During the pilot, 1,192 smart bands and 317 millimeter-wave sleep-radar units were deployed. Across the 10-month period, 4.286 million raw records were collected. Records fell into three source groups (Table 1).

Data tier	Examples	Update frequency	Privacy risk
Terminal real-time sensing	Heart rate / SpO ₂ , sleep radar activity,	Per minute / per event	High for location; low for physiology

	gas/door/window sensors		
Primary-care clinical	Chronic-disease follow-up, annual check-up, prescriptions, diagnosed history	Per visit	High (legally sensitive)
Community service	Online consults, offline care work orders, family interactions	Per service event	Moderate (re-identification risk when aggregated)

Table 1. Three groups of heterogeneous data sources on the Hui-YangLao platform.

3.3 Framework overview

As illustrated in Figure 1, the proposed framework has four layers. The bottom layer ingests the three heterogeneous source types. The second layer applies tiered desensitization according to privacy risk. The third layer runs a three-stage lightweight fusion pipeline. The top layer exposes differently granular data views to four classes of users. A quarterly feedback loop tunes perturbation strength and feature weights between measured risk-identification accuracy and privacy-alert volume.

Figure 1. Tiered desensitization and three-stage lightweight fusion architecture deployed on the "Hui-YangLao" platform




Figure 1. Tiered desensitization and three-stage lightweight fusion architecture deployed on the Hui-YangLao platform.

4. Tiered Desensitization Scheme

4.1 Risk-tier classification

Following the principle that protection strength should be proportionate to disclosure risk [15,16], we assigned each source attribute to one of three tiers. High-tier attributes include precise building/unit location, severe-disease diagnoses, and direct identifiers (name, phone number, national ID). Medium-tier

attributes include chronic-disease follow-up, routine check-up results, and prescriptions. Low-tier attributes include routine vital signs (heart rate, sleep duration) and aggregated work-order statistics.

4.2 Per-tier protection operations

High-tier. Precise building and unit numbers were permanently deleted and replaced with a fuzzy community-area code. Severe-disease and severe-disability records were retained locally at the community health center and not synchronized to the cloud. Direct identifiers were replaced by irreversible pseudonyms, and no plaintext identifier was transmitted over wide-area networks.

Medium-tier. Disease and drug names were mapped to a non-invertible code table; the cloud stores only the coded values. Full EMRs remain at the clinic; clinicians request access through an online workflow requiring dual approval.

Low-tier. Routine vital signs were protected with a small additive perturbation ($\pm 3\%$) that preserves monthly and annual trend curves. Device and work-order identifiers were protected with a lightweight hash. All desensitization operations were executed on dual nodes — the terminal IoT gateway and the platform data middle-end — so that plaintext records never traverse long-haul networks.

5. Three-stage Lightweight Fusion Pipeline

5.1 Stage 1 — Data-level cleaning

Heterogeneous sources used different timestamps, identifier schemes, and update rates. Stage 1 unified timestamps to a common clock, mapped each older adult to a stable virtual identifier, removed sensor outliers (e.g., impossible heart-rate values or static-radar false events), and imputed short follow-up gaps. Of the 4.286 million raw records, 0.314 million (7.32%) were removed as noise, leaving 3.972 million cleaned records.

5.2 Stage 2 — Feature-level weighted fusion

Three feature groups were constructed — chronic-disease indicators, home-activity indicators, and environmental indicators — with chronic features assigned higher weights than daily-activity features. The weighted operation produced a static health profile vector; the raw source records were not retained after this stage. Per-batch server CPU/memory consumption was measured against the platform's previous unfused baseline.

5.3 Stage 3 — Decision-level cross fusion

In Stage 3, the health feature vector, home-activity sequence, and care work orders were cross-validated to issue a small set of graded risk labels — fall, chronic worsening, and home-safety — rather than pushing raw vital signals or location. Only these labels reached downstream dashboards, blocking sensitive raw data at the output boundary.

5.4 Four-tier role-based access control

Four roles were defined. Level 1 (emergency): temporary fuzzy-location retrieval during rescue, with every retrieval logged. Level 2 (clinicians): coded chronic-disease datasets after approval. Level 3 (community grid staff): aggregate district-level risk statistics only. Level 4 (platform operations / third-party vendors): aggregated statistical indicators only, with no access to individual records. All permission changes and ad-hoc retrievals were retained for at least three years to support regulatory audit.

6. Evaluation

6.1 Metrics and measurement protocol

We evaluated the deployment along three axes. (i) Privacy: a composite identity re-identification risk score on a 0–10 scale, produced by the platform's internal audit module using a deterministic linkage check against quasi-identifiers (age, gender, community area, rare disease codes); and the geolocation re-identification probability, estimated by simulating an attacker who holds the released dataset and a public

resident list. (ii) Data utility: R^2 between original and desensitized chronic-disease trend series; chronic-risk identification accuracy, measured against clinician-labeled ground truth on a held-out validation subset; mean absolute error of heart-rate readings; and curve-matching R^2 for monthly sleep duration. (iii) Operational: per-batch CPU/memory overhead compared with the pre-pipeline baseline; mean approval latency for EMR retrieval; count of unauthorized-access alerts; and volume of operation logs. Results are reported as platform-monitored values; no external red-team evaluation was conducted, which we acknowledge as a limitation.

6.2 Privacy protection outcomes

After high-tier desensitization, the composite identity re-identification risk score dropped from 9.2/10 to 1.1/10 (an 88.0% relative reduction). Geolocation re-identification probability fell from 100% to 0.70%. At the medium tier, the share of sensitive fields in cloud-stored records dropped from 41.7% to 0%. Across the 10-month pilot, 17 high-risk access events (bulk export, off-hours login, or repeated queries) were flagged and closed within 24 hours; no confirmed exfiltration of identifiable elderly health records was recorded by the platform's internal monitoring.

6.3 Data utility outcomes

Utility loss was small. Chronic-trend fitting R^2 decreased from 0.961 to 0.948 (a 1.35% relative drop). Chronic-risk identification accuracy decreased from 92.36% to 90.82% — a 1.54-percentage-point change. At the low tier, mean absolute heart-rate error was 1.82 beats/minute; monthly sleep-duration curve matching R^2 was 0.983; and aggregate work-order statistics error was below 0.4%. The final fused risk labels maintained a stable identification accuracy of at least 90.15% across the pilot.

Metric	Before desensitization	After desensitization	Change
Identity re-ID risk score (0–10)	9.2	1.1	–88.0%
Geolocation re-ID probability	100%	0.70%	–99.3 pp
Cloud sensitive-field ratio	41.7%	0%	–41.7 pp
Chronic-trend fitting R^2	0.961	0.948	–0.013
Chronic-risk identification accuracy	92.36%	90.82%	–1.54 pp

Table 2. Privacy and utility metrics before and after tiered desensitization.

6.4 Computational overhead and governance

Feature-level weighted fusion reduced per-batch computational overhead by 61.7% relative to the previous raw-data fusion baseline, allowing the pipeline to run on a low-cost community cloud server. EMR retrieval approval averaged 4.7 minutes. The monitoring module logged 124,000 operation events over the pilot. Six quarterly parameter iterations were performed: when warning false/negative rates rose, perturbation strength on the affected tier was relaxed; when unauthorized-access alerts rose, protection was tightened. Over these iterations, risk-identification accuracy improved by 4.21 percentage points.

7. Discussion

The deployment makes three concrete contributions to the literature. First, it shows that proportionate, tiered protection — rather than uniform encryption or blanket anonymization — is a practical middle path for community elderly-care data. The 1.54-percentage-point drop in chronic-risk accuracy and the 0.013 drop in trend-fitting R^2 are small compared with the 88% reduction in re-identification risk. This pattern is consistent with reports that well-configured de-identification can preserve downstream utility [17,15], but it

is demonstrated here on a real-world community deployment rather than a benchmark dataset.

Second, the three-stage fusion pipeline is deliberately organized so that the most expensive operations (feature weighting and decision cross-validation) run only on compressed vectors, while raw multi-source records never leave the immediate care environment. The 61.7% reduction in per-batch overhead is operationally significant: it is what makes the entire scheme affordable for a county or sub-district platform. This responds directly to Gravina et al.'s [21] call for hardware-aware fusion design and to the observation that most state-of-the-art elderly monitoring methods assume hospital-grade compute [26,22].

Third, the scheme pairs technical controls with governance: a four-role access model, dual-approval EMR retrieval, and a quarterly feedback loop between accuracy and privacy alerts. This aligns with Sun et al.'s [10] finding that trust and human oversight, rather than technology alone, determine whether older adults accept smart-care systems. The 17 closed-loop alerts and the absence of confirmed exfiltration events are encouraging, although — as we discuss below — these are internal platform metrics and do not substitute for external audit.

Compared with recent international work, our approach is not algorithmically novel: we use established pseudonymization, generalization, light perturbation, and multi-level fusion rather than proposing a new DP mechanism. Its value lies in the operational configuration — which tier gets which mechanism, at what privacy budget, on what hardware, and with what governance — for a setting that is largely absent from the algorithmic literature. We expect the configuration to be adaptable to other mid-sized community platforms in China and comparable middle-income urban contexts.

8. Limitations

Several limitations should be noted. First, the sample is drawn from six urban communities and three townships in a single district of Chengdu; findings may not transfer unchanged to rural China, to rapidly aging megacities, or to other national contexts. Second, privacy metrics come from the platform's internal audit module rather than an external red-team evaluation; an independent re-identification study would strengthen the evidence. Third, the accuracy figures rely on clinician-labeled ground truth on a held-out subset; inter-rater reliability was not formally reported here. Fourth, although written informed consent was obtained, we did not measure older adults' actual privacy preferences or trust perceptions, which other studies have shown to be heterogeneous across disease types and data recipients [38]. Fifth, the current "full-process traceability" rests on operation logs that could in principle be altered by privileged insiders; we plan to explore permissioned distributed-ledger storage in future work. Finally, the 10-month window is sufficient for feasibility assessment but too short to evaluate long-term outcomes such as hospitalization reduction or mortality.

9. Conclusion

For resource-constrained community elderly-care platforms, balancing privacy compliance, computational cost, and personalized health-service accuracy is a persistent practical problem. Based on a 10-month deployment on the Hui-YangLao platform, we have described a tiered desensitization scheme coupled with a lightweight three-stage fusion pipeline and a four-role access-control model. Measured reductions in re-identification risk were large (identity risk score 9.2 to 1.1; geolocation re-identification 100% to 0.70%), while utility loss remained modest (chronic-risk accuracy 92.36% to 90.82%; trend R^2 0.961 to 0.948) and computational overhead dropped by 61.7%. The configuration offers a low-cost, reproducible reference implementation that other community platforms can adapt to local hardware, regulatory, and stakeholder conditions. Future work will extend the evaluation to rural terminals, integrate tamper-resistant audit logging, and add formal measurement of older adults' privacy preferences and long-term health outcomes.

Declarations

Ethics approval and consent to participate: Ethics approval was obtained from the [Institutional Review

Board name, approval number to be inserted] prior to deployment. Written informed consent was obtained from each participant or their legal representative. Consent for publication: Not applicable. All data were anonymized; no individual details, images, or videos are included. Availability of data and materials: The datasets analyzed during the current study are not publicly available due to the sensitive nature of older adults' health records and applicable personal information protection regulations, but are available from the corresponding author on reasonable request and subject to institutional data-sharing approval.

References

- [1] Sweeney, L. (2002). k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(5), 557–570.
- [2] Machanavajjhala, A., Kifer, D., Gehrke, J., & Venkitasubramaniam, M. (2006). l-diversity: Privacy beyond k-anonymity. *Proceedings of the 22nd International Conference on Data Engineering (ICDE)*.
- [3] Dwork, C. (2006). Differential privacy. *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP)*, LNCS 4052, 1–12.
- [4] Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4), 211–407.
- [5] National Bureau of Statistics of China. (2024). 2023 national economic and social development statistical bulletin. Beijing: NBS.
- [6] Zhang, Q., Li, M., & Wu, Y. (2020). Smart home for elderly care: development and challenges in China. *BMC Geriatrics*, 20, 317.
- [7] Wang, B., & Xu, L. (2021). Construction of the "Internet Plus" community smart elderly care service platform. *Journal of Healthcare Engineering*, 2021, 4310648.
- [8] Hanpo, H., He, W., Yingchao, W., Jin, Y., & Guixing, Q. (2022). Construction path of smart health care platform for the elderly in China. *Strategic Study of Chinese Academy of Engineering*, 24(2).
- [9] Wei, Y., Liu, X., & Hou, R. (2023). User profile in smart elderly care community: Findings from a community in western China. *Journal of Beijing Institute of Technology*, 32(2).
- [10] Sun, Y., Ankenbauer, S. A., Guo, Z., Chen, Y., Ma, X., & He, L. (2025). Rethinking technological solutions for community-based older adult care: Insights from "Older Partners" in China. *Proceedings of the ACM on Human-Computer Interaction*.
- [11] Psychoula, I., Merdivan, E., Singh, D., Chen, L., Chen, F., Hanke, S., Kropf, J., Holzinger, A., & Geist, M. (2018). A deep learning approach for privacy preservation in assisted living. *arXiv:1802.09359*.
- [12] Jourdan, T., Boutet, A., Bahi, A., & Frindel, C. (2021). Privacy-preserving IoT framework for activity recognition in personal healthcare monitoring. *Proceedings on Privacy Enhancing Technologies*.
- [13] Chowdhury, M. J. M., Kayes, A. S. M., Watters, P., Scoler-Gray, P., Ng, A., & Dillon, T. (2020). Patient controlled, privacy preserving IoT healthcare data sharing framework. *Proceedings of the 53rd Hawaii International Conference on System Sciences (HICSS)*.
- [14] Mohammed, N., Jiang, X., Chen, R., Fung, B. C. M., & Ohno-Machado, L. (2013). Privacy-preserving heterogeneous health data sharing. *Journal of the American Medical Informatics Association (JAMIA)*, 20(3), 462–469.
- [15] Bonomi, L., Gousheh, S., & Fan, L. (2023). Enabling health data sharing with fine-grained privacy. *Proceedings on Privacy Enhancing Technologies*.
- [16] Li, X., Zhao, H., Yu, D., Wang, L., & Liu, P. (2019). Multidimensional correlation hierarchical differential privacy for medical data with multiple privacy requirements. *Proceedings of the 2nd International Conference on Healthcare Science and Engineering, LNEE 536*.
- [17] Soria-Comas, J., Domingo-Ferrer, J., Sánchez, D., & Martínez, S. (2014). Enhancing data utility in differential privacy via microaggregation-based k-anonymity. *The VLDB Journal*, 23(5), 771–794.

- [18] Adams, T., Birkenbihl, C., Otte, K., Ng, H. G., Rieling, J. A., Ahrer, A. N., Sax, U., Prasser, F., & Fröhlich, H. (2025). On the fidelity versus privacy and utility trade-off of synthetic patient data. *iScience*.
- [19] Almasi, M. M., Siddiqui, T. R., Mohammed, N., & Hemmati, H. (2016). The risk-utility tradeoff for data privacy models. *Proceedings of the IEEE International Conference on Data Science and Advanced Analytics*.
- [20] Hassan, M. U., Rehmani, M. H., & Chen, J. (2019). Differential privacy techniques for cyber physical systems: A survey. *IEEE Communications Surveys & Tutorials*.
- [21] Gravina, R., Alinia, P., Ghasemzadeh, H., & Fortino, G. (2017). Multi-sensor fusion in body sensor networks: State-of-the-art and research challenges. *Information Fusion*, 35, 68–85.
- [22] Wang, X., Ellul, J., & Azzopardi, G. (2020). Elderly fall detection systems: A literature survey. *Frontiers in Robotics and AI*, 7, 71.
- [23] Muhammad, G., Alshehri, F., Karray, F., El Saddik, A., Alsulaiman, M., & Falk, T. H. (2022). A comprehensive survey on multimodal medical signals fusion for smart healthcare systems. *Information Fusion*, 76, 375–395.
- [24] Nweke, H. F., Wah, T. Y., Mujtaba, G., & Al-Garadi, M. A. (2019). Data fusion and multiple classifier systems for human activity detection and health monitoring: Review and open research directions. *Information Fusion*, 46, 147–170.
- [25] John, A., Cardiff, B., & John, D. (2024). A review on multisensor data fusion for wearable health monitoring. *arXiv:2412.05895*.
- [26] Sun, H., & Chen, Y. (2023). A rapid response system for elderly safety monitoring using progressive hierarchical action recognition. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*.
- [27] Baumgartner, M., Kreiner, K., Lauschensky, A., et al. (2024). Health data space nodes for privacy-preserving linkage of medical data to support collaborative secondary analyses. *Frontiers in Medicine*, 11, 1301660.
- [28] Lin, X., Gu, X., Liang, M., Liao, T., Liang, H., & Yu, X. (2025). Key technologies and application of data-asset-element-driven integrated medical-nursing-rehabilitation care. *China Digital Medicine*, 20(4), 14–20.
- [29] Jiang, S., Liu, Y., Li, Q., Tao, W., Zhu, B., & Xing, P. (2023). Research progress on human fall detection based on multi-source heterogeneous data fusion. *Chinese Journal of Gerontology*, 43(4), 1010–1013.
- [30] Gu, L., Fang, N., & Song, Y. (2025). Mechanism and optimization paths of digital technology enabling value co-creation in smart elderly-care services: A case study of L sub-district, Shanghai. *Urban Problems*, (3), 28–38.
- [31] Zuo, M. (2014). Connotation, models and opportunities of smart elderly care. *China Public Security*, (10), 48–50.
- [32] Zuo, M., Duan, R., & Lin, L. (2024). Research on smart elderly-care data resource governance system: An activity-theory perspective. *Journal of Xi'an Jiaotong University (Social Sciences)*, 44(3), 40–51.
- [33] Su, D. (2025). Smart healthy elderly-care models and protection of older adults' privacy information. *Human Rights Law Review*, (1), 113–135.
- [34] Wang, L., & Hao, X. (2023). Research progress and prospects on supply–demand matching of medical-nursing integration services in China. *Chinese Health Policy Research*, 16(2), 1–8.
- [35] Dong, Y., & Liu, S. (2025). Public data authorized operation as a solution to home-based smart elderly-care dilemmas. *Journal of Sichuan Normal University (Social Sciences)*, 52(5), 112–121.
- [36] Tang, J. (2024). The "care–caring" theory, the elderly care system and a care-oriented society.

Social Policy Research, (1), 64–71.

[37] Sun, Y., & Li, X. (2025). A hypergraph convolution-based intelligent healthcare platform for aging population management. *Frontiers in Public Health*, 13, 1640808.

[38] Calero Valdez, A., & Ziefle, M. (2018). The users' perspective on the privacy–utility trade-offs in health recommender systems. *International Journal of Human-Computer Studies*.