

# Multimodal Sensor Fusion for Robust Physiological Monitoring during Free-Living Activities

Dongjie Huangfu

Department of Computer Science, George Mason University, Fairfax, VA, USA.  
huangfu-dongjie@gmu.edu

Kasper C. Watson

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.  
kasper-c-watson@buffalo.edu

Emmett L. Benson

Department of Computer Science, University of North Texas, Denton, TX, USA.  
emmett.b@unt.edu

Samir Garrett

Department of Computer Science, University of Central Florida, Orlando, FL, USA.  
samirmail@ucf.edu

## Abstract

The expansion of wearable sensing technologies has created new opportunities for continuous physiological monitoring beyond controlled clinical environments. However, free-living conditions introduce substantial challenges related to motion artifacts, contextual variability, heterogeneous sensor placements, environmental interference, and population diversity. Single-modality physiological sensing often fails to provide sufficiently robust estimates during unstructured daily activities. Multimodal sensor fusion offers a promising system-level strategy by combining complementary physiological, kinematic, and environmental data streams to improve estimation reliability, contextual awareness, and clinical interpretability. This paper presents a long-form analysis of multimodal sensor fusion for physiological monitoring during free-living activities. It examines the architectural dimensions of sensing, edge processing, communication, and cloud analytics while emphasizing structural trade-offs among accuracy, energy consumption, latency, privacy, and scalability. The discussion explores early, intermediate, and decision-level fusion approaches and their implications for system robustness. It further considers context-aware adaptation, motion artifact mitigation, data governance, algorithmic fairness, population heterogeneity, and deployment sustainability. Rather than focusing narrowly on signal-level algorithms, the paper adopts a systems-oriented perspective, integrating concerns from artificial intelligence, wearable engineering, socio-technical infrastructure, and health policy. The analysis highlights the need for governance frameworks that address continuous physiological data collection and emphasizes the importance of fairness across demographic groups. Forward-looking perspectives on edge intelligence, federated learning, digital biomarkers, and sustainable wearable ecosystems are provided. The paper argues that robust free-living physiological monitoring requires not only advanced fusion algorithms but also carefully designed sociotechnical architectures that align technical performance with ethical and regulatory requirements.

## **Keywords**

physiological monitoring, multimodal sensor fusion, wearable systems, free-living activities, robustness, data governance, fairness, edge computing.

## **1. Introduction**

The increasing availability of wearable devices capable of measuring photoplethysmography, electrocardiography, electrodermal activity, skin temperature, and movement has transformed the possibilities for continuous health assessment. Wearable sensors are now used for remote health monitoring in applications ranging from cardiovascular surveillance to stress estimation and sleep analysis [1]. However, moving physiological monitoring from controlled laboratories into free-living environments introduces considerable complexity. In free-living settings, users engage in diverse physical activities, experience changing environmental conditions, and interact with systems in ways that produce substantial variability in signal quality and physiological interpretation. Activity detection and classification using wearable sensors have therefore become essential components of contextual physiological monitoring [2]. For example, photoplethysmographic signals are highly susceptible to motion artifacts during intensive physical exercise, which can severely distort heart rate estimation [3]. These limitations motivate the development of multimodal systems that combine multiple sensor streams to compensate for the weaknesses of individual modalities.

Multimodal sensor fusion has emerged as a foundational approach for improving the reliability of wearable physiological monitoring. Multi-sensor fusion in body sensor networks has been studied extensively to reduce uncertainty and increase robustness in dynamic environments [4]. By combining kinematic data from accelerometers and gyroscopes with physiological signals, a system can identify activity contexts, detect artifact-prone periods, and adapt its processing strategies accordingly. Advances in multimodal deep learning further enable systems to learn joint representations from heterogeneous data sources, capturing interactions that are difficult to engineer manually [5]. These developments suggest that robust physiological monitoring is not merely a signal processing problem but a systems integration challenge involving sensor design, feature extraction, machine learning architectures, and real-world deployment constraints.

Despite the promise of multimodal fusion, significant gaps remain in translating laboratory performance into free-living reliability. The healthcare internet of things has expanded the connectivity of wearable devices, yet many systems still struggle with energy constraints, data quality variability, and interoperability [6]. Moreover, physiological monitoring systems must address governance and fairness concerns that become more acute when data are collected continuously from diverse populations. This paper examines multimodal sensor fusion from a systems perspective, emphasizing architectural decisions, structural trade-offs, deployment sustainability, robustness, fairness, and policy implications. The goal is to provide an interdisciplinary synthesis that bridges engineering design and socio-technical governance for free-living physiological monitoring.

## **2. Related Work and System Context**

Research on wearable health monitoring has progressed from single-sensor proof-of-concept studies to integrated systems that incorporate multiple modalities, edge computing, and distributed learning. The healthcare internet of things has been widely surveyed as a framework for connecting sensors, gateways, and cloud services to support remote care [6]. Within this ecosystem, physiological data are generated continuously, creating demands for

real-time processing and reliable communication. However, centralized cloud processing introduces latency and privacy risks, especially when sensing occurs in homes, workplaces, and public spaces. Federated learning has been proposed as a privacy-preserving method for training machine learning models across distributed devices without transferring raw physiological data to central servers [7]. This approach aligns with the need to learn from diverse free-living data while limiting exposure of sensitive health information.

Motion artifact reduction remains one of the most active research areas in wearable photoplethysmography and physiological sensing. Multiwavelength approaches have been explored to improve heart rate estimation from wrist-worn sensors by leveraging differences in tissue absorption and motion sensitivity across optical channels [8]. In parallel, multi-sensor data fusion techniques have been developed to integrate accelerometry, physiological signals, and environmental context for more robust biomedical inference [9]. These methods range from traditional ensemble classifiers to deep neural architectures that learn hierarchical representations across modalities. Edge cognitive computing has also been introduced to reduce dependence on cloud resources by enabling local inference and adaptive sensor management in smart healthcare systems [10]. Together, these lines of work form the technical context for multimodal fusion in free-living physiological monitoring.

The system context is further shaped by the need to move beyond algorithmic accuracy toward operational viability. Free-living monitoring requires systems to function over extended periods with limited battery capacity, intermittent connectivity, and highly variable user behaviors. The structural choices made at the architecture level determine whether a fusion algorithm can be deployed effectively outside controlled settings. Consequently, the present discussion integrates related work on signal fusion with broader considerations of edge infrastructure, privacy, fairness, and sustainability.

### **3. Architectural Overview of Multimodal Physiological Monitoring**

A robust multimodal physiological monitoring system can be conceptualized as a layered architecture comprising sensing, on-device processing, communication, data aggregation, and higher-level analytics. The sensing layer includes heterogeneous transducers that capture physiological and contextual information, such as optical photoplethysmography sensors, electrode-based biopotential sensors, inertial measurement units, and environmental sensors. Each modality has distinct physical principles, noise profiles, bandwidth requirements, and energy costs. Multi-sensor fusion architectures must therefore manage heterogeneity at the lowest level while preserving information that may be relevant for higher-level interpretation [4].

The on-device processing layer performs tasks such as filtering, artifact detection, feature extraction, and local inference. Edge processing is particularly important in free-living monitoring because it reduces the volume of raw data that must be transmitted and enables rapid adaptation to changing contexts. However, edge devices impose strict computational and energy constraints, leading to structural trade-offs between model complexity and inference latency. Multimodal deep learning models can achieve powerful representational fusion, but their computational demands may exceed the capacity of low-power wearable processors [5]. Systems designers must therefore decide which portions of the fusion pipeline are executed locally and which are offloaded to more capable devices or cloud services.

The communication layer determines how data move between wearable nodes, personal gateways, and remote infrastructure. Wireless protocols such as Bluetooth Low Energy and

Wi-Fi offer different balances of energy consumption, bandwidth, range, and reliability. Free-living conditions introduce variability in connectivity that may require buffering, compression, or opportunistic transmission. Healthcare internet of things architectures address these issues through layered protocols and data management strategies, but many deployments remain fragmented [6]. The cloud analytics layer provides long-term storage, retrospective model training, and population-level analysis. Yet overreliance on cloud infrastructure can undermine real-time responsiveness and create concentrated privacy risks. Federated approaches can distribute learning across devices, allowing model improvement without centralizing raw physiological data [7]. These architectural choices directly affect system robustness, scalability, and governance.

System-level design must also account for the temporal dynamics of free-living behavior. Physiological signals are nonstationary and context-dependent, and multimodal fusion systems must adapt their processing to both short-term activity bursts and longer-term physiological trends. An effective architecture therefore includes context detection, state estimation, and feedback loops that allow sensing parameters and inference models to be adjusted dynamically. Such adaptability improves robustness but introduces complexity in verification and validation, as system behavior may change during deployment.

#### **4. Sensor Fusion Strategies and Structural Trade-offs**

Multimodal sensor fusion can be organized along several conceptual levels. Early fusion combines raw or lightly processed signals before feature extraction, allowing the model to learn joint representations from the full data stream. This approach can capture subtle cross-modal interactions but may be sensitive to noise and missing data. Feature-level fusion combines modality-specific features into a shared representation, preserving domain-specific preprocessing while enabling joint inference. Decision-level fusion aggregates outputs from modality-specific classifiers, offering robustness to individual sensor failure but potentially losing fine-grained correlations. In body sensor networks, these strategies have been compared in terms of accuracy, energy consumption, and scalability [4].

The choice of fusion level has implications for system architecture. Early fusion often requires synchronized high-bandwidth signal streams, which increases data transmission and processing demands. Feature-level fusion can reduce communication load if features are computed locally, but it requires careful feature design and may introduce information loss. Decision-level fusion is more modular and fault-tolerant, yet it may fail to exploit complementary information when single-modality classifications are uncertain. Multimodal deep learning methods can learn hierarchical representations that combine these levels in flexible ways, but they also complicate deployment on constrained hardware [5]. Structural trade-offs therefore emerge between representational power, computational complexity, and operational robustness.

Beyond accuracy, multimodal systems must address broader performance attributes including fairness, interpretability, and regulatory compliance. Machine learning models trained on biased data can produce systematic disparities in health-related predictions. Algorithmic bias in healthcare has been documented in multiple contexts, demonstrating that apparently accurate systems may still fail underserved populations [11]. Fairness considerations must be integrated into the fusion design process rather than treated as post-hoc checks [12]. Deep learning models in healthcare can provide substantial benefits but also raise challenges related to transparency, validation, and accountability [13]. Digital health systems face ethical and policy pressures related to consent, data ownership, and the social implications of continuous

monitoring [14]. These factors influence architectural choices such as the granularity of data storage, the availability of audit logs, and the use of interpretable fusion components.

Another structural trade-off concerns the distribution of intelligence across edge and cloud tiers. Edge-based fusion reduces latency and privacy exposure but limits the complexity of models that can be deployed. Cloud-based fusion supports sophisticated learning and population-scale analytics but introduces dependency on connectivity and central infrastructure. Hybrid architectures attempt to balance these concerns by running lightweight adaptive models on-device while periodically synchronizing with cloud services for model updates and retrospective analysis. The structural choices made here influence not only system performance but also the governance and ownership of physiological data.

## **5. Robustness and Context-Aware Adaptation in Free-Living Conditions**

Free-living physiological monitoring must contend with a broad range of disturbances that are rarely present in controlled studies. Motion artifacts, sensor displacement, changes in ambient light, perspiration, and variable skin contact all degrade signal quality. Activity intensity varies continuously, from sedentary periods to vigorous exercise, and the corresponding physiological responses differ in amplitude and time course. Rapid vital sign extraction methods have been developed to support real-time opto-physiological monitoring across varying physical activity intensity levels [15]. Such methods use adaptive signal processing and motion reference signals to estimate physiological parameters more reliably during dynamic conditions.

Context-aware adaptation is essential for robust free-living performance. Activity recognition from inertial sensors can inform the fusion system about the user's current movement state, enabling it to adjust filtering strategies, sensor weighting, or model selection [2]. For example, during walking or running, accelerometer readings can serve as artifact references for photoplethysmography denoising [3]. The fusion system can also use activity context to interpret physiological changes, distinguishing between heart rate elevation caused by physical exertion and elevation associated with stress or illness. This contextual interpretation improves both clinical relevance and false alarm reduction.

Despite these advances, robustness remains difficult to guarantee across diverse user populations and environments. Optical heart rate sensors have been shown to exhibit inaccuracies related to skin tone, sensor fit, and motion patterns, highlighting the need for population-specific validation and algorithmic adjustment [16]. Free-living studies also face challenges related to adherence and data completeness, as users may remove devices or engage in activities that are not well represented in training data [17]. Multimodal fusion can mitigate some of these issues by providing redundancy and cross-modal validation, but it cannot fully eliminate biases introduced during data collection and model training.

A system-level approach to robustness therefore includes sensor redundancy, context-aware processing, continuous quality assessment, and adaptive model updating. The fusion architecture should monitor the reliability of each modality in real time and adjust its confidence accordingly. If one sensor becomes unreliable, the system can rely more heavily on others while flagging periods of reduced certainty. This dynamic weighting strategy improves resilience but requires careful calibration to avoid feedback loops that reinforce poor estimates. Moreover, robust systems must be evaluated in uncontrolled settings, where ground truth is often noisy or incomplete. Validation frameworks must therefore move

beyond aggregate accuracy metrics to include per-context performance, failure mode analysis, and subgroup-specific assessments.

## **6. Data Governance, Privacy, and Ethical Implications**

Continuous physiological monitoring generates highly sensitive data that can reveal not only health status but also behavioral patterns, emotional states, and daily routines. The governance of such data involves multiple layers of responsibility, including device manufacturers, application developers, cloud providers, healthcare institutions, and regulatory bodies. Federated learning offers one mechanism for reducing privacy risks by keeping raw data on devices and sharing only model updates [7]. However, federated systems are not immune to privacy attacks, and additional safeguards such as differential privacy, secure aggregation, and access controls are needed. The implementation of these safeguards must be integrated into the fusion architecture rather than added as an afterthought.

Consent and transparency are particularly challenging in free-living monitoring because data collection is continuous, passive, and often context-dependent. Users may not fully understand how their physiological data are combined with location, activity, and environmental information to produce health inferences. Digital health frameworks must address the ethical implications of long-term sensing and the potential for function creep, where data collected for one purpose are later used for another [14]. Governance structures should include clear data purpose limitations, user-facing controls, and mechanisms for data deletion or correction.

The multimodal nature of the data also raises identifiability concerns. Even if individual modalities are de-identified, the combination of activity patterns, physiological responses, and environmental markers may create unique behavioral fingerprints. The fusion process itself can therefore amplify privacy risks by linking otherwise innocuous data streams into a more revealing profile. System designers must consider data minimization, local processing, and the use of aggregated or abstracted features to limit exposure. At the same time, excessive abstraction can reduce clinical utility, creating a structural tension between privacy and diagnostic value.

Regulatory oversight must evolve alongside technical capabilities. Existing frameworks for medical devices and data protection may not fully address adaptive fusion systems that update their models over time. Continuous monitoring systems can change their behavior after deployment, complicating pre-market review and post-market surveillance. Governance models should incorporate algorithmic impact assessments, independent audits, and reporting requirements for performance disparities across populations. These mechanisms can help ensure that multimodal fusion systems are not only accurate but also accountable.

## **7. Fairness and Population Heterogeneity**

Fairness in physiological monitoring requires attention to differences in sensor performance, model generalization, and health outcomes across demographic groups. Biased training data can lead to systematic errors in heart rate estimation, stress detection, or activity classification. For example, photoplethysmography signals can vary with skin pigmentation, and models trained predominantly on lighter skin tones may perform poorly for darker skin tones [16]. If such models are deployed without subgroup validation, they may produce inequitable health information. Algorithmic fairness research has shown that health algorithms can encode and amplify existing disparities when training data reflect structural inequities [11]. Fairness is

therefore not solely a property of the dataset but also of the system design, evaluation procedures, and deployment context [12].

Population heterogeneity extends beyond skin tone to include age, body composition, cardiovascular fitness, sex, and the presence of chronic conditions. These factors influence sensor contact, signal morphology, and physiological baseline values. A fusion system trained on young, healthy adults may fail to generalize to older adults with different vascular properties or to patients with arrhythmias. Multimodal systems must incorporate population-specific priors and allow for individualized calibration without sacrificing scalability. Adaptive personalization can improve fairness by adjusting thresholds and model weights to match individual characteristics, but it also introduces new governance challenges related to explainability and consistency.

Fairness assessments should be conducted across relevant subgroups and across the range of activities encountered in free-living settings. Evaluation protocols should include diverse participant recruitment, transparent subgroup reporting, and analysis of worst-case performance. The activity tracker literature demonstrates that intervention effects and measurement accuracy can vary substantially across populations and study designs, underscoring the importance of robust validation beyond convenience samples [17]. In practice, achieving fairness requires collaboration among sensor engineers, clinical researchers, data scientists, and community stakeholders to define meaningful performance criteria and to detect unintended harms.

## **8. Deployment, Sustainability, and Infrastructure Considerations**

The deployment of multimodal physiological monitoring systems at scale involves substantial infrastructure investment and long-term sustainability planning. Wearable devices must be manufactured, distributed, maintained, and eventually recycled or disposed of. Soft and skin-mounted sensor systems have demonstrated the potential for high-quality physiological sensing outside laboratory settings, but their production costs and durability remain concerns [18]. The materials used in flexible electronics, batteries, and adhesive components have environmental implications that are often overlooked in proof-of-concept research. Sustainable deployment requires attention to device lifetime, repairability, and end-of-life management alongside technical performance.

Data infrastructure is another key consideration. Continuous multimodal sensing produces large volumes of data that must be stored, processed, and analyzed over time. The promise of big data in healthcare depends on the ability to turn raw physiological streams into clinically meaningful insights while managing computational and storage costs [19]. However, large-scale data centralization also increases security risks and governance burdens. Distributed architectures that combine edge processing with selective cloud aggregation can reduce infrastructure demands and improve resilience. The optimal balance depends on the intended application, the regulatory environment, and the available connectivity.

Clinical integration remains a significant barrier to deployment. For wearable systems to influence care, they must fit into existing workflows, electronic health record systems, and reimbursement structures. The emerging clinical role of wearables is shaped by factors such as evidence quality, interoperability, clinician trust, and patient engagement [20]. Multimodal fusion systems must therefore be designed not only for technical accuracy but also for usability by clinicians and patients. This requires standardized outputs, clear uncertainty

communication, and interfaces that support clinical decision-making without adding cognitive burden.

Finally, long-term sustainability depends on building institutional capacity for continuous monitoring programs. This includes training healthcare professionals, establishing data governance committees, and developing policies for responsible use. Interdisciplinary collaboration is essential, as no single engineering or clinical domain can address the full range of challenges. The long-term value of multimodal physiological monitoring will be realized only if systems are deployable, maintainable, fair, and aligned with societal values.

## **9. Conclusion**

Multimodal sensor fusion represents a critical pathway toward robust physiological monitoring in free-living environments. By combining complementary sensing modalities, systems can overcome the limitations of individual sensors, adapt to dynamic activity contexts, and provide more reliable health-related inferences. However, the technical promise of fusion cannot be realized without careful attention to architectural design, structural trade-offs, governance, fairness, and deployment sustainability. This paper has examined these dimensions from a systems perspective, emphasizing the need to integrate signal-level innovation with broader socio-technical considerations.

The architectural choices made in sensing, edge processing, communication, and cloud analytics determine not only accuracy but also privacy, energy consumption, and scalability. Fusion strategies such as early, feature-level, and decision-level approaches offer different balances of representational power and operational feasibility. Context-aware adaptation and rapid vital sign extraction methods improve robustness during varying physical activity intensity, yet population-specific validation remains essential. Data governance frameworks must address the unique risks of continuous multimodal monitoring, including identifiability, consent, and algorithmic bias. Fairness across demographic groups must be treated as a central design objective rather than a secondary concern.

Future work should focus on developing sustainable edge intelligence, federated personalization, and transparent validation frameworks for free-living monitoring. Interdisciplinary collaboration among engineers, clinicians, ethicists, and policymakers will be essential to translate multimodal fusion research into equitable and effective health systems. By integrating technical performance with societal responsibility, next-generation physiological monitoring systems can achieve robust, trustworthy, and inclusive health monitoring beyond the clinic.

## **References**

1. Majumder, S., Mondal, T., & Deen, M. J. (2017). Wearable sensors for remote health monitoring. *Sensors*, 17(1), 130.
2. Cornacchia, M., Ozcan, K., Zheng, Y., & Velipasalar, S. (2017). A survey on activity detection and classification using wearable sensors. *IEEE Sensors Journal*, 17(2), 386-403.
3. Zhang, Z., Pi, Z., & Liu, B. (2015). TROIKA: A general framework for heart rate monitoring using wrist-type photoplethysmographic signals during intensive physical exercise. *IEEE Transactions on Biomedical Engineering*, 62(2), 522-531.

4. 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, 68-80.
5. Ngiam, J., Khosla, A., Kim, M., Nam, J., Lee, H., & Ng, A. Y. (2011). Multimodal deep learning. *Proceedings of the 28th International Conference on Machine Learning*, 689-696.
6. Islam, S. M. R., Kwak, D., Kabir, M. H., Hossain, M., & Kwak, K.-S. (2015). The internet of things for health care: A comprehensive survey. *IEEE Access*, 3, 678-708.
7. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., ... Cardoso, M. J. (2020). The future of digital health with federated learning. *NPJ Digital Medicine*, 3, 119.
8. Zhang, Y., Song, S., Vullings, R., Biswas, D., Simões-Capela, N., van Helleputte, N., van Hoof, C., & Groenendaal, W. (2019). Motion artifact reduction for wrist-worn photoplethysmograph sensors based on different wavelengths. *Sensors*, 19(3), 673.
9. Muzammal, M., Talat, R., Sodhro, A. H., & Pirbhulal, S. (2020). A multi-sensor data fusion enabled ensemble approach for medical data from body sensor networks. *Information Fusion*, 53, 155-164.
10. Chen, M., Li, W., Hao, Y., Qian, Y., & Humar, I. (2018). Edge cognitive computing based smart healthcare system. *Future Generation Computer Systems*, 86, 403-411.
11. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447-453.
12. 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), 1-35.
13. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., Cui, C., Corrado, G., Thrun, S., & Dean, J. (2019). A guide to deep learning in healthcare. *Nature Medicine*, 25(1), 24-29.
14. Vayena, E., Haeusermann, T., Adjekum, A., & Blasimme, A. (2018). Digital health: Meeting the ethical and policy challenges. *Swiss Medical Weekly*, 148, w14571.
15. Zheng, X., Dwyer, V. M., Barrett, L. A., Derakhshani, M., & Hu, S. (2023). Rapid vital sign extraction for real-time opto-physiological monitoring at varying physical activity intensity levels. *IEEE Journal of Biomedical and Health Informatics*, 27(7), 3107-3118.
16. Bent, B., Goldstein, B. A., Kibbe, W. A., & Dunn, J. P. (2020). Investigating sources of inaccuracy in wearable optical heart rate sensors. *NPJ Digital Medicine*, 3, 18.
17. Laranjo, L., Ding, D., Heleno, B., Kocaballi, B., Quiroz, J. C., Tong, H. L., ... Lau, A. Y. S. (2021). Do smartphone applications and activity trackers increase physical activity in adults? Systematic review, meta-analysis and metaregression. *British Journal of Sports Medicine*, 55(8), 422-432.
18. Xu, S., Zhang, Y., Jia, L., Mathewson, K. E., Jang, K.-I., Kim, J., ... Rogers, J. A. (2014). Soft microfluidic assemblies of sensors, circuits, and radios for the skin. *Science*, 344(6179), 70-74.
19. Shilo, S., Rossman, H., & Segal, E. (2020). Axes of a revolution: Challenges and promises of big data in healthcare. *Nature Medicine*, 26(1), 29-38.

20. Smuck, M., Odonkor, C. A., Wilt, J. K., Schmidt, N., & Swiernik, M. A. (2021). The emerging clinical role of wearables: Factors for successful implementation in healthcare. *NPJ Digital Medicine*, 4, 45.