

GAN-Enhanced Reconstruction of Motion-Corrupted Physiological Signals for Wearable Health Monitoring

Richard M. Jacobs

School of Electrical Engineering and Computer Science, Oregon State University, Corvallis,
OR, USA.

jacobs1978@oregonstate.edu

Claude Schmidt

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

s.claude@binghamton.edu

Abstract

Wearable health monitoring systems increasingly collect photoplethysmographic, electrocardiographic, and bioimpedance signals during everyday activity, but ambulatory motion introduces complex and nonstationary artifacts that degrade clinical and wellness inferences. Conventional filtering and handcrafted signal decomposition methods often fail when noise overlaps with physiological bandwidths or when motion patterns vary across individuals and contexts. This paper examines the system-level role of generative adversarial networks in reconstructing motion-corrupted physiological signals, focusing not on a single algorithmic formulation but on architectural trade-offs, infrastructure constraints, governance requirements, and deployment pathways. The paper discusses how adversarial training can learn a manifold of plausible physiological waveforms and produce reconstructions that preserve diagnostically relevant morphology while suppressing motion-related distortion. The analysis addresses conditional architectures, domain adaptation, sensor heterogeneity, and the need for lightweight models suitable for edge execution. It further considers fairness risks arising from training data that underrepresent skin tones, body compositions, and motion profiles, as well as the sustainability implications of training and deploying deep generative models in wearable ecosystems. The paper argues that GAN-enhanced reconstruction should be treated as a socio-technical infrastructure rather than an isolated signal processing module, requiring careful attention to validation, regulatory oversight, privacy, and long-term maintenance. A forward-looking perspective is offered for responsible integration into remote monitoring, digital clinical trials, and population health platforms.

Keywords

wearable health monitoring; generative adversarial networks; motion artifact correction; physiological signal reconstruction; responsible deployment; edge computing.

1. Introduction

Wearable physiological sensors have moved from controlled laboratory settings into continuous everyday use, generating longitudinal data streams that support cardiovascular screening, sleep staging, stress estimation, and management of chronic conditions. Photoplethysmography, electrocardiography, electrodermal activity, and bioimpedance provide noninvasive windows into autonomic and circulatory function. However, the

interpretive value of these signals is tightly coupled to their quality under free-living conditions. Motion corruption remains one of the most persistent barriers because daily activities produce accelerations, skin stretch, sensor displacement, and microvascular changes that overlap spectrally and temporally with the physiological signal of interest. The resulting degradation cannot be fully resolved by frequency-domain filtering, since motion artifacts often occupy the same frequency band as heart rate or respiratory components.

Generative adversarial networks have emerged as a powerful class of deep generative models for synthesizing and reconstructing high-dimensional data [1]. In the physiological monitoring context, a generator can be trained to transform corrupted segments into cleaner waveforms that resemble an implicitly learned distribution of valid physiological signals, while a discriminator attempts to distinguish reconstructed examples from clean references [2]. This adversarial formulation encourages reconstructions that are difficult to separate from real physiological waveforms, moving beyond pointwise error minimization toward structural plausibility. Wasserstein formulations have been developed to stabilize training and provide more meaningful convergence signals [3], and improved training procedures further reduce instability in adversarial optimization [4]. Deep convolutional architectures have also provided useful inductive biases for learning local temporal features in generative models [5]. These advances form the basis for reconstructing motion-corrupted physiological waveforms in wearable systems.

This paper adopts a system-level perspective on GAN-enhanced reconstruction for wearable health monitoring. Rather than presenting a single algorithm or benchmark, it examines the structural design choices, governance mechanisms, fairness implications, and infrastructure constraints that determine whether such models can be responsibly deployed outside controlled evaluations. The paper is organized as follows. Section 2 introduces the system context and background. Section 3 analyzes motion artifact corruption in wearable signals. Section 4 discusses GAN-based reconstruction principles. Section 5 examines architectural considerations and structural trade-offs. Section 6 addresses data governance, bias, and fairness. Section 7 explores deployment, sustainability, and infrastructure. Section 8 considers policy and clinical translation. Section 9 concludes.

2. Background and System Context

A wearable health monitoring system typically comprises sensing hardware, local preprocessing, wireless communication, cloud or edge analytics, and an interface for clinicians or users. Each layer introduces distinct constraints on signal reconstruction. For example, a wrist-worn optical sensor may have limited power, memory, and computational resources, while a cloud-based service can host larger models but must contend with latency, privacy, and connectivity variability. In this layered context, motion artifact correction is not only a signal processing problem but also a system design problem. A reconstruction model that achieves high accuracy on a laboratory benchmark may be unsuitable for edge deployment if its inference latency drains the battery or if its memory footprint exceeds the microcontroller capacity.

Early wearable heart rate monitoring approaches often combined accelerometer references with spectral subtraction or adaptive filtering [6]. Time-varying spectral filtering methods later provided improved robustness by tracking the signal and noise characteristics over short windows [7]. Reviews of wearable photoplethysmography have emphasized that while hardware improvements reduce some artifacts, algorithmic resilience remains essential for clinical-grade monitoring during ambulatory activity [8]. The transition from handcrafted

signal models to learned representations has been gradual but significant, because learned models can capture nonlinear relationships between auxiliary motion signals and physiological waveform distortions that fixed mathematical transformations cannot easily express.

Generative adversarial networks introduced a framework in which two neural networks compete: one generates candidates, and the other distinguishes real from generated examples [1]. Conditional variants allow external information, such as corrupted input segments or accelerometer signals, to guide generation [2]. Wasserstein formulations and their improved training procedures stabilize adversarial learning and reduce mode collapse, making them more applicable to one-dimensional physiological signals where training instability can be particularly harmful [3], [4]. Deep convolutional architectures have further provided useful inductive biases for learning local temporal features [5]. These advances form the basis for reconstructing motion-corrupted physiological waveforms in wearable systems.

3. Motion Artifact Corruption in Wearable Physiological Monitoring

Motion artifacts in wearable physiological signals arise through several interacting mechanisms. Physical acceleration displaces the sensor relative to the skin, altering optical coupling in photoplethysmography or electrode impedance in electrocardiography. Tissue deformation changes blood volume distribution and scattering, introducing baseline wander and abrupt amplitude changes. Ambient light interference and muscle activation add further distortion. Because these processes are nonstationary and context dependent, the artifact characteristics differ across walking, running, typing, sleeping, and postural transitions. A signal segment recorded during rhythmic running may exhibit quasi-periodic artifact patterns, while a sudden arm movement may produce transient high-amplitude spikes.

The central difficulty is that motion artifacts do not simply add independent noise to a clean physiological signal. They modulate the signal path, obscure waveform landmarks, and can mimic arrhythmias or other clinically significant events. Traditional filtering methods often rely on assumptions about signal and noise frequency bands, but the overlap between typical heart rate frequencies and motion frequencies means that aggressive filtering can remove real physiological information. Reconstruction is therefore not only denoising but also inference about the underlying physiological state from heavily distorted observations. In wearable photoplethysmography, for example, motion can cause peak detection errors that cascade into incorrect heart rate variability estimates [7].

Deep learning for time-series data has shown that learned hierarchical representations can be more robust to such structured corruption than fixed feature extractors [9]. However, a major challenge is that the space of valid physiological waveforms is highly constrained by anatomy, cardiovascular dynamics, and sensor placement, while the space of motion corruption is broad and diverse. A generative approach can exploit this asymmetry by learning a manifold of plausible clean physiological signals and using it to constrain reconstruction. This motivates the use of generative adversarial networks, which provide a mechanism for matching the distribution of reconstructed signals to the distribution of clean references rather than merely minimizing average error across a training set.

4. Generative Adversarial Networks for Signal Reconstruction

In GAN-enhanced reconstruction, the generator receives a motion-corrupted physiological segment and, optionally, auxiliary signals such as accelerometer readings or skin temperature, and produces a corrected version. The discriminator receives pairs of input and output signals

and learns to determine whether the output is a real clean reference or a generator-produced reconstruction. Through iterative training, the generator improves its ability to produce outputs that the discriminator cannot reliably distinguish from clean signals. This adversarial process can preserve fine-grained waveform morphology that might be lost by mean squared error optimization, which tends to produce overly smooth reconstructions.

Conditional generative adversarial networks are particularly relevant because they allow the reconstruction to be conditioned on the corrupted input and contextual sensor data [2]. The discriminator can also evaluate the consistency between the corrupted input and the reconstructed output, discouraging the generator from inventing physiologically implausible features. Wasserstein generative adversarial networks replace the original objective with a formulation that provides more stable gradients and meaningful convergence signals [3]. Improved training techniques, such as gradient penalty, further stabilize the optimization and are important for physiological time series where training collapse could yield clinically dangerous hallucinations [4]. Deep convolutional generators and discriminators can learn multiscale temporal structures, from local beat morphology to longer respiratory modulation [5].

Despite these advantages, GAN-based reconstruction introduces structural risks. The generator may remove subtle but diagnostically important abnormalities if those abnormalities are underrepresented in the clean training distribution. It may also generate plausible but false waveform features under severe corruption. Therefore, the adversarial objective must be combined with task-specific validation and uncertainty estimation. A system that uses GAN reconstruction as a preprocessing stage should propagate reconstruction uncertainty to downstream algorithms, such as heart rate estimation, arrhythmia detection, or oxygen saturation inference. The clinical acceptability of the reconstruction depends not only on visual similarity but also on the preservation of features that downstream decision models use.

5. Architectural Considerations and Structural Trade-offs

The design of a GAN-enhanced physiological signal reconstruction system involves multiple architectural decisions. A one-dimensional convolutional generator can efficiently capture local temporal patterns and is well suited to wearable signals with limited input lengths. Recurrent or attention-based components can model longer dependencies, but they increase memory and computational cost. Hybrid architectures that combine convolutional encoders with attention mechanisms can balance local feature extraction and global context. The choice of generator and discriminator capacities must reflect the target hardware, the latency budget, and the complexity of the motion artifact distribution. Overly large models may overfit to the training environment and fail to generalize to new sensor placements or activity types.

Domain shift is a major concern in wearable deployment. A model trained on data from young adults performing controlled treadmill protocols may perform poorly for older patients with different gait patterns or for individuals with darker skin tones, higher body mass, or peripheral vascular conditions. Adversarial domain adaptation methods can align feature distributions across source and target domains [10], [11]. A reconstruction model may use adversarial training not only to distinguish clean from corrupted signals but also to make the learned representation invariant to sensor type, body site, or demographic variables. However, such invariance must be designed carefully, because physiological differences that are clinically meaningful should not be erased.

Attention mechanisms have been explored for opto-physiological monitoring during physical activity to prioritize relevant signal regions while suppressing motion-dominated segments [12]. This line of research suggests that attention can improve reconstruction quality by adaptively weighting temporal features, but it also introduces additional computational overhead and may require careful regularization to prevent attention from focusing on spurious correlations. Model compression techniques such as pruning, quantization, and Huffman coding have been developed to reduce the memory and computational footprint of deep networks without severely degrading accuracy [13]. In wearable systems, compressed generators can be executed on microcontrollers or low-power application processors, but aggressive compression may reduce the model's ability to represent rare physiological events, creating a structural trade-off between efficiency and clinical sensitivity. Another structural trade-off concerns the granularity of reconstruction. A system may reconstruct entire waveform segments, individual beats, or only latent features needed for downstream tasks. Reconstructing the full waveform preserves morphology and supports clinical inspection, but it requires more computation and storage. Reconstructing only derived features is more efficient but risks losing information that may be needed for future analyses or for auditability. The choice should be guided by the intended use case, the available hardware, and the governance requirements for transparency and traceability.

Domain adaptation remains critical for wearable deployment because sensor placement, skin properties, and activity patterns differ across populations. Adversarial domain adaptation methods can align feature distributions across source and target domains [10], [11]. However, alignment should not remove clinically meaningful variation. For example, differences in pulse wave morphology between older adults and younger adults may reflect arterial stiffness rather than sensor artifact. A reconstruction model that treats all demographic variation as nuisance could systematically distort diagnostic information. Governance structures should therefore require subgroup-specific validation and fairness audits. A modular system architecture in which the reconstruction model is separated from downstream diagnostic models can facilitate such audits by making the reconstruction step independently testable.

6. Data Governance, Bias, and Fairness

The performance of GAN-enhanced physiological signal reconstruction depends heavily on the composition of the training data. If the clean reference signals are collected predominantly from individuals with light skin tones, low body mass, and young adult cardiovascular systems, the learned manifold of plausible physiological waveforms will be skewed toward those characteristics. Motion artifacts in photoplethysmography are known to vary with melanin concentration because optical absorption differs across skin tones, and the same sensor may produce weaker or noisier signals for some populations. A reconstruction model trained on homogeneous data may fail to produce clean waveforms for underrepresented groups, or it may generate waveforms that do not preserve their physiological features. This becomes a fairness problem when the reconstructed signals are used for downstream decisions such as arrhythmia screening or oxygen saturation estimation. Bias can therefore be propagated from the training data through the generative model into clinical or wellness recommendations.

Data governance for wearable health monitoring must account for the entire lifecycle of training and deployment. Raw physiological signals are sensitive health data, and centralized collection for model training raises privacy risks. Federated learning offers an alternative in which model updates, rather than raw signals, are shared across devices [14]. This approach

can reduce the movement of sensitive data and enable learning from diverse populations without centralizing identifiable records. However, federated learning introduces its own governance challenges, including heterogeneity across devices, the risk of reconstruction attacks on shared gradients, and the need for robust aggregation protocols. Differential privacy mechanisms can provide formal bounds on the leakage of individual information during training [15]. When combined with federated learning, differential privacy can support a stronger privacy posture, but the added noise may degrade the quality of the reconstructed waveforms, particularly for rare physiological events. System designers must balance privacy protection against clinical sensitivity and be transparent about these trade-offs.

Fairness cannot be reduced to a single metric. A GAN-based reconstruction model might show similar average signal-to-noise improvement across demographic groups while still producing disproportionate errors in specific subgroups during high-motion conditions. Intersectional analysis is needed, considering combinations of skin tone, age, body composition, sensor placement, and activity profile. Audit frameworks should evaluate not only overall reconstruction error but also the preservation of clinically meaningful waveform features across subgroups. In addition, the quality of the reference labels themselves must be questioned. Clean reference signals are often obtained in controlled settings that do not reflect free-living conditions, and label noise can introduce systematic biases into the generative model. Transparent documentation of data sources, collection protocols, and population characteristics is essential for meaningful oversight.

7. Deployment, Sustainability, and Infrastructure

Deploying GAN-enhanced reconstruction in wearable ecosystems requires careful consideration of where inference occurs. Cloud-based processing offers abundant computation and easy model updates, but it depends on continuous connectivity and may not meet latency requirements for real-time monitoring. Edge inference preserves local data and reduces communication overhead, but it imposes strict limits on model size, memory, and energy consumption. A hybrid architecture may be appropriate, in which lightweight models perform initial reconstruction on the device and more complex models refine ambiguous segments in the cloud. The decision about where to run the model is not merely technical; it affects privacy, reliability, and the sustainability of the overall system. Wearable devices have limited battery capacity, and energy-intensive inference can shorten device lifetime or increase charging frequency, which in turn affects user adherence.

Explainability and uncertainty estimation are particularly important for generative models in health monitoring. A reconstruction model may produce a waveform that appears plausible but is not grounded in the observed physiological signal. Techniques for visualizing which parts of the input influenced the output can help developers and regulators understand model behavior [16]. Uncertainty estimation can inform downstream systems when a reconstruction is reliable and when it should be treated with caution. In clinical contexts, high-performance artificial intelligence systems must earn trust through transparency and consistent performance across diverse settings [17]. The integration of GAN-based reconstruction into clinical workflows should therefore be accompanied by mechanisms for reviewing reconstruction quality and for flagging segments that may contain hallucinations.

Regulatory frameworks for software as a medical device are evolving to address adaptive and data-driven technologies. The United States Food and Drug Administration has proposed an action plan for artificial intelligence and machine learning based medical software, emphasizing the importance of predetermined change control plans, real-world performance

monitoring, and transparency [18]. A GAN-enhanced reconstruction module that serves as a preprocessing step for diagnostic algorithms may be subject to regulatory oversight if its outputs influence clinical decisions. Manufacturers should plan for post-market surveillance and for updates that may alter model behavior. Federated learning environments for digital health further complicate regulatory review because models trained across distributed devices may evolve without a single centralized artifact [19]. Clear documentation of model versions, training data distributions, and performance characteristics is necessary for accountability.

Sustainability is an often overlooked dimension of deep generative model deployment. Training and inference for large neural networks consume substantial computational resources and energy [20]. In wearable health monitoring, the cumulative energy footprint of millions of devices running continuous inference can be significant. Model compression, quantization, and efficient architecture design can reduce per-inference energy use, but they may also reduce model capacity. Life-cycle assessments should consider not only the energy used during training but also the energy consumed by on-device inference, data transmission, and model updates. A responsible deployment strategy should avoid unnecessary complexity and align model capacity with actual clinical or wellness needs. The goal is not maximal accuracy in isolation but sustainable and equitable benefit across populations and over time.

8. Policy and Clinical Translation

The translation of GAN-enhanced reconstruction from research prototypes to clinical practice depends on more than technical performance. Validation studies must demonstrate that reconstructed signals preserve the features required for specific clinical tasks, such as detecting atrial fibrillation, estimating respiratory rate, or tracking sleep quality. Benchmarks based solely on waveform similarity may not capture clinical utility. Studies should include diverse patient populations, realistic motion protocols, and comparisons with established signal processing methods. The absence of standardized evaluation protocols for generative models in physiological monitoring makes it difficult to compare results across studies and to establish the safety of deployed systems. Professional societies, regulators, and standards bodies could contribute by defining reference datasets, performance metrics, and reporting requirements.

Interoperability is another policy concern. Wearable devices from different manufacturers use different sensors, sampling rates, and data formats. A reconstruction model trained for one device may not transfer to another without adaptation. Standardized data formats and application programming interfaces can reduce friction and enable third-party validation. However, interoperability must be balanced against the need to protect proprietary algorithms and business interests. Policy instruments may encourage the development of open benchmarks and reference implementations without mandating disclosure of all model details. The goal should be to create an ecosystem in which independent evaluation is possible and in which users and clinicians can understand the limitations of the algorithms they rely on.

Liability and accountability become complex when a generative model alters physiological signals before clinical interpretation. If a missed arrhythmia can be traced to a reconstruction error, responsibility may be distributed among the device manufacturer, the algorithm developer, the healthcare provider, and the platform operator. Clear documentation of algorithmic limitations and the circumstances under which reconstruction should not be trusted is essential. Clinical decision support systems should not present reconstructed waveforms as raw data without indication of the processing applied. Transparency about the generative nature of the reconstruction can help clinicians interpret the information

appropriately. In addition, patients should be informed about how their data are processed and about the potential for algorithmic error.

Equity considerations extend across geographic and economic boundaries. High-quality wearable monitoring and advanced reconstruction algorithms may be deployed first in wealthy urban populations, while communities with limited connectivity or older devices may not benefit. Global health applications often require robustness to low-cost sensors, intermittent connectivity, and extreme environmental conditions. Generative models trained on high-end device data may perform poorly in these settings. Policy frameworks should support technology transfer, open datasets that include diverse global populations, and capacity building in low-resource regions. The societal benefit of GAN-enhanced physiological monitoring will be limited if it widens existing digital health divides.

9. Conclusion

GAN-enhanced reconstruction of motion-corrupted physiological signals offers a promising pathway toward more reliable wearable health monitoring in free-living conditions. By learning a manifold of plausible physiological waveforms, adversarial training can suppress motion artifacts while preserving fine-grained morphology that may be lost by conventional filtering. However, the system-level challenges are substantial. Architectural choices must balance reconstruction fidelity against computational cost, latency, and energy consumption. Domain adaptation must preserve clinically meaningful variation rather than erase demographic differences. Data governance must address privacy, consent, and bias, while deployment strategies must account for regulatory oversight, interoperability, and sustainability. The use of attention mechanisms within generative models can improve robustness during physical activity, but such mechanisms must be implemented with care to avoid spurious correlations and excessive resource demands [12]. A responsible approach treats GAN-enhanced reconstruction as a socio-technical infrastructure embedded within larger monitoring and decision-making systems. Its success depends not only on adversarial optimization but also on transparent validation, equitable data practices, and governance structures that align technical capability with public trust and clinical benefit.

References

1. Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27, 2672–2680.
2. Mirza, M., & Osindero, S. (2014). Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*.
3. Arjovsky, M., Chintala, S., & Bottou, L. (2017). Wasserstein GAN. *arXiv preprint arXiv:1701.07875*.
4. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., & Courville, A. (2017). Improved training of Wasserstein GANs. *Advances in Neural Information Processing Systems*, 30, 5767–5777.
5. Radford, A., Metz, L., & Chintala, S. (2016). Unsupervised representation learning with deep convolutional generative adversarial networks. In *Proceedings of the International Conference on Learning Representations*.

6. 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.
7. Peng, F., Zhang, Z., Gou, X., Liu, H., & Wang, W. (2014). Motion artifact removal from photoplethysmographic signals by combining temporally constrained independent component analysis and adaptive filter. *Biomedical Engineering Online*, 13, 50.
8. Allen, J. (2007). Photoplethysmography and its application in clinical physiological measurement. *Physiological Measurement*, 28(3), R1–R39.
9. Ismail Fawaz, H., Forestier, G., Weber, J., Idoumghar, L., & Muller, P.-A. (2019). Deep learning for time series classification: a review. *Data Mining and Knowledge Discovery*, 33(4), 917–963.
10. Ganin, Y., & Lempitsky, V. (2015). Unsupervised domain adaptation by backpropagation. In *Proceedings of the 32nd International Conference on Machine Learning*, 1180–1189.
11. Tzeng, E., Hoffman, J., Saenko, K., & Darrell, T. (2017). Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7167–7176.
12. Zheng, X., Hu, S., Dwyer, V., Barrett, L., & Derakhshani, M. (2026). Joint attention mechanism learning to facilitate opto-physiological monitoring during physical activity. *Biomedical Signal Processing and Control*, 113, 108949.
13. Han, S., Mao, H., & Dally, W. J. (2016). Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In *Proceedings of the International Conference on Learning Representations*.
14. McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 1273–1282.
15. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318.
16. Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision*, 618–626.
17. Topol, E. J. (2019). High-performance medicine: the convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44–56.
18. U.S. Food and Drug Administration. (2021). Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan. U.S. Department of Health and Human Services.
19. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier-Hein, K., Ourselin, S., Sheller, M., Summers, R. M., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. *npj Digital Medicine*, 3, 119.

20. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 3645–3650.