# Adaptive State Space Networks for Robust Multi-View Analysis of Complex Medical Imaging Data

Pavel D. James
Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
paveldj@buffalo.edu

**Abstract**

Complex medical imaging datasets increasingly require integrated analysis across multiple views and modalities, including computed tomography, magnetic resonance imaging, positron emission tomography, ultrasound, and histopathology. These views differ in resolution, contrast, anatomical orientation, and noise characteristics, but collectively provide complementary diagnostic evidence. Adaptive state space networks have emerged as a promising architectural framework for multi-view medical imaging because they combine efficient long-range sequence modeling with input-dependent state transitions. However, clinical translation requires attention to systemic factors beyond predictive accuracy, including infrastructure constraints, robustness under missing data, interpretability, fairness, regulatory accountability, and environmental sustainability. This paper presents a systems-oriented analysis of adaptive state space networks for robust multi-view medical image analysis. It examines architectural foundations and structural trade-offs associated with cross-view state propagation, missing-view handling, and fusion granularity. It further discusses deployment considerations related to hospital information systems, distributed training, federated learning, latency, and hardware requirements. The paper also addresses robustness and uncertainty estimation, interpretability of state trajectories, demographic fairness, governance mechanisms, and regulatory implications. Sustainability is analyzed in terms of model compression, energy consumption, and lifecycle management. By emphasizing the interplay between algorithmic design and sociotechnical infrastructure, the paper argues that adaptive state space models can support clinically meaningful multi-view analysis only when architectural choices are aligned with institutional workflows, equity objectives, and monitoring systems. Future directions include multimodal integration, causal reasoning, realistic benchmarking, and governance frameworks for continuous learning systems.

**Keywords**

adaptive state space networks, multi-view medical imaging, robust representation learning, clinical deployment, fairness, governance, sustainability.

## 1. Introduction

Modern medical imaging produces complex and heterogeneous data across multiple views that capture different anatomical, physiological, and pathological characteristics. A single patient may undergo several imaging studies, each generated by a distinct modality with its own spatial resolution, contrast mechanism, and diagnostic purpose. The integration of these views can improve diagnostic confidence, treatment planning, and longitudinal monitoring, but it also introduces significant computational and organizational challenges. Recent advances in selective state space sequence modeling [1] and their visual extensions [2] provide a new pathway for processing long image sequences without the quadratic complexity associated with dense attention. These models are particularly relevant for multi-view analysis

because they can selectively retain and propagate information across extended spatial contexts. Nevertheless, their introduction into clinical imaging environments requires careful consideration of architecture, infrastructure, and governance rather than an exclusive focus on benchmark performance.

The development of medical image analysis has been strongly influenced by convolutional neural networks. Architectures such as U-Net [3] and later self-configuring pipelines such as nnU-Net [4] established reliable baselines for segmentation and image-level prediction. Attention-based transformers [5], vision transformers [6], and hierarchical windowed models [7] expanded the capacity to capture global relationships. Residual and dense connectivity strategies [8], [9] likewise contributed to optimization stability and feature reuse. These advances have enabled impressive results in many tasks, but multi-view imaging introduces problems of alignment, asynchronous acquisition, and heterogeneous feature spaces that are not fully addressed by single-stream architectures. Adaptive state space networks can potentially address these problems by modeling each view as a sequence and learning cross-view state interactions.

Clinical applications of deep learning in dermatology and chest radiography have demonstrated that models can approach specialist-level performance in selected tasks [10], [11]. However, these studies often use curated datasets and single-view inputs. Multi-view diagnostic systems require more complex explanatory tools because decisions may depend on evidence distributed across several images. Post hoc interpretation methods such as gradient-weighted class activation mapping [12] can highlight relevant regions in one image, but they do not fully reveal how views influence one another. The central argument of this paper is that adaptive state space networks should be designed as sociotechnical components rather than isolated predictive models. Their clinical value depends on robustness, transparency, resource efficiency, and alignment with institutional governance structures. The following sections develop this argument through architectural analysis and systems engineering perspectives.

## 2. Background and Related Work

Medical image analysis has progressed through several architectural paradigms, each introducing distinct inductive biases. Convolutional networks remain important because they encode spatial locality and translation equivariance, which are useful for identifying local anatomical structures. U-Net [3] introduced symmetric encoder-decoder pathways with skip connections that preserve fine spatial detail, while nnU-Net [4] demonstrated that systematic configuration choices can produce robust segmentation performance across many datasets. Residual networks [8] address optimization difficulties in deep models, and densely connected networks [9] improve feature reuse and parameter efficiency. These convolutional foundations continue to influence modern medical imaging pipelines, especially when computational resources are limited or when local features are highly diagnostic.

Attention-based models introduced a more flexible approach to global context modeling. The transformer architecture [5] enables interactions across all tokens, while vision transformers [6] apply this principle to image patches. Hierarchical models such as Swin Transformer [7] reduce computational cost by limiting attention to local windows while still enabling cross-window communication. These models have achieved strong results in medical imaging tasks, but their full attention mechanisms can be expensive for high-resolution volumetric data. When multiple views are combined, the sequence length grows rapidly, and quadratic attention becomes a critical bottleneck. This limitation motivates interest in selective state

space models that maintain a compact latent state while processing long sequences [1]. Visual state space models [2] adapt this idea to image inputs by using bidirectional scanning paths and patch-based sequence construction.

Multi-view fusion has historically been categorized into early, late, and intermediate strategies. Early fusion integrates images or features before prediction, which can be effective when views are well aligned but is vulnerable to missing data. Late fusion combines decisions from separate models, offering robustness to missing views but potentially missing cross-view interactions. Intermediate fusion exchanges learned features during processing, allowing more flexible interaction but requiring careful design. Recent cross-fusion state space models for multi-view medical image classification [13] show that structured cross-view scanning can improve classification by allowing one view to modulate the processing of another. This direction suggests a broader shift from static fusion policies toward adaptive mechanisms that decide dynamically how much information should be shared. However, clinical translation requires more than improved classification accuracy; it demands robust behavior under incomplete data, clear interpretability, and compatibility with healthcare infrastructure.

## 3. Architectural Foundations of Adaptive State Space Networks for Multi-View Imaging

Adaptive state space networks extend traditional sequence models by making state transitions dependent on the input. In a conventional state space formulation, recurrent updates follow fixed rules that may not distinguish between relevant and irrelevant information. Selective mechanisms modify this behavior by allowing the model to update its internal state based on the current input, effectively deciding what to remember and what to discard. In multi-view imaging, this selective updating can be applied at two levels. First, each image can be divided into patches and processed as a sequence, with the scanning order determining how spatial context is accumulated. Second, sequences from different views can be interleaved or processed through cross-view state transitions. This nested structure allows the model to build representations that are sensitive to both local anatomy and cross-modal relationships.

A central architectural consideration is the granularity of fusion. View-specific encoders can project heterogeneous inputs into a shared latent space, after which shared adaptive state space layers perform cross-view reasoning. This separation preserves modality-specific feature extraction while enabling common representations for diagnosis, prognosis, or segmentation. The approach can be extended with bidirectional scanning so that information flows both forward and backward across a view and across views. Such bidirectional cross-scanning can improve classification [13], but it also creates dependencies between views. If one view is missing, the model must still produce a coherent representation. Architectural solutions include learned placeholders, view-dropout regularization during training, and state initialization strategies that allow inference under partial observation. These mechanisms help ensure that the model degrades gracefully rather than failing unpredictably when data are incomplete.

From a systems perspective, adaptive state space networks are not merely computational graphs. They are embedded within pipelines that include image acquisition, quality control, registration, harmonization, and clinical interpretation. Patch size, scanning order, feature dimensionality, and cross-view routing all influence memory usage, latency, and hardware utilization. State space models offer favorable asymptotic complexity for long sequences, but actual deployment performance depends on implementation details such as parallel scan algorithms and memory access patterns. The architecture should therefore be designed with awareness of downstream infrastructure. Modular view encoders can simplify updates when

new imaging devices or protocols are introduced. A tightly coupled model may achieve slightly better internal alignment but can be harder to maintain under evolving clinical conditions. The choice between modularity and tight coupling is a structural trade-off with long-term operational consequences.

## 4. System-Level Trade-Offs and Infrastructure Design

Deploying adaptive state space networks in clinical environments requires integration with picture archiving and communication systems, electronic health records, and clinical decision support interfaces. A deployed model must ingest images from these systems, apply preprocessing, perform inference, and return results within clinically acceptable timeframes. Although state space models can process long sequences efficiently, sequential state updates may underutilize parallel hardware unless specialized scan implementations are used. System designers must evaluate whether inference can be performed on hospital data center accelerators or on edge devices closer to the point of care. Latency requirements differ across tasks. Emergency settings demand near real-time results, while retrospective analysis can tolerate longer processing times. The infrastructure must also support logging, versioning, and audit trails so that model outputs can be investigated if clinical concerns arise.

Training multi-view state space networks introduces additional complexity. Medical image datasets are large, high dimensional, and often distributed across institutions. Distributed training across multiple accelerators may be necessary, but state space models have communication patterns that may differ from convolutional or transformer architectures. Pipeline parallelism, gradient accumulation, and mixed-precision training can improve scalability, yet they increase system complexity and operational risk. Federated learning is an attractive approach for preserving privacy while training across hospitals, but it raises issues related to data heterogeneity, communication efficiency, and security. Cross-view fusion modules in a federated setting may require aggregation strategies that respect institutional differences in imaging protocols. If one site contributes only certain views, the global model must be robust to partial participation and non-identically distributed data. These considerations are as important as raw algorithmic performance.

Fusion granularity also affects infrastructure design. Early fusion is computationally simple but sensitive to registration and missing views. Late fusion is modular and robust but may lose important cross-view interactions. Adaptive intermediate fusion seeks to combine the benefits of both by learning dynamic information exchange. This approach increases routing complexity and may complicate quality assurance because the influence of each view is less straightforward. Nevertheless, adaptive intermediate fusion is appealing because it resembles clinical reasoning, in which specialists integrate evidence across modalities while maintaining separate uncertainty about each source. To support this design, systems should record state trajectories and fusion decisions where feasible, enabling retrospective review. Without such infrastructure, clinicians may be unable to determine why a decision changed when a particular view became available or was removed.

Interoperability is another critical requirement. Models must exchange data with heterogeneous hospital systems, which often use different formats, identifiers, and data quality conventions. Standardized interfaces and metadata schemas can reduce integration costs. Model outputs should include confidence estimates, input metadata, and version identifiers. These requirements are not secondary to architecture; they shape the design of the entire system. A state space model that is highly accurate in a controlled research environment may fail in practice if its inputs are inconsistently preprocessed or if its outputs

cannot be combined with existing clinical workflows. Therefore, systems engineering and model architecture should evolve together, with explicit attention to failure modes and operational constraints.

**5. Robustness, Uncertainty,5. Robustness, Uncertainty, and Interpretability**

Robustness in multi-view medical image analysis extends beyond accuracy on fixed test sets to include the capacity of a model to behave safely under distribution shift, missing views, heterogeneous acquisition protocols, and adversarial perturbations. Adaptive state space networks can contribute to robustness through their selective state updates, which in principle allow the model to ignore irrelevant or noisy inputs while preserving diagnostically important evidence. However, this capability is not automatic. The state transition mechanism must be trained under conditions that reflect realistic clinical variability. A model that performs well when all views are present may fail when a view is unavailable because of equipment malfunction, patient motion, or scheduling constraints. Robustness therefore requires explicit simulation of missing data during training and evaluation. View dropout, where entire views are randomly omitted, encourages the network to avoid over-reliance on any single modality. This form of regularization can be combined with learned placeholder states so that the model maintains coherent state trajectories even when input sequences are incomplete. Without such techniques, cross-view fusion mechanisms may create brittle dependencies that undermine diagnostic reliability.

Uncertainty estimation is closely related to robustness. Clinical decision-making requires not only a prediction but also an indication of confidence. Bayesian deep learning and ensemble methods provide established approaches for estimating predictive uncertainty in medical imaging [14]. State space networks can be adapted to produce uncertainty estimates by maintaining stochastic state representations or by training ensembles with different initialization and data orderings. Monte Carlo dropout applied to selective state transitions is possible, but the latent state dynamics introduce additional considerations because uncertainty may accumulate across long sequences. Calibration is particularly important when a model integrates multiple views, since the contribution of each view may vary with image quality and patient condition. A well-calibrated model should express high uncertainty when views are conflicting or degraded. This requires evaluation protocols that measure calibration under view deletion, noise injection, and domain shift, rather than relying solely on area under the receiver operating characteristic curve. Clinical acceptance depends on whether uncertainty information can be communicated in a way that supports, rather than disrupts, reasoning processes.

Interpretability of multi-view state space networks poses a distinct set of challenges. Post hoc saliency methods such as gradient-weighted class activation mappings [12] can identify spatial regions that influence a decision, but they often treat the model as a static function and may not reveal how evidence from one view modulates the interpretation of another. State space models introduce an additional layer of complexity because the internal state evolves over the sequence. It may be possible to visualize state trajectories and quantify how strongly each view changes the latent representation. Such trajectory-based explanations could be valuable in clinical settings, where specialists need to understand whether a diagnosis depends primarily on one high-quality view or on converging evidence from multiple sources. Nevertheless, the latent state is not necessarily aligned with human-understandable concepts, and trajectory plots may be misleading if not evaluated rigorously. Interpretability research should therefore move toward concept-level explanations that connect state transitions to

clinically meaningful features such as lesion boundaries, tissue texture, and anatomical landmarks. These explanatory tools must be developed in collaboration with clinicians to ensure that they address actual decision-making needs.

## 6. Fairness, Bias, and Governance

The deployment of multi-view medical imaging models raises serious fairness concerns because training data may reflect historical inequalities in healthcare access, diagnostic practices, and population representation. Algorithms trained on datasets dominated by certain demographic groups may perform less accurately for underrepresented groups, a problem that has been documented in widely used health prediction systems [18]. In medical imaging, bias can arise not only from demographic imbalances but also from differences in imaging equipment, referral patterns, and annotation quality. Multi-view systems may amplify or obscure these biases depending on how they aggregate evidence. If a particular view is more frequently available or of higher quality for one population, the model may learn to rely on that view and underperform when it is missing or degraded for another population. Fairness therefore requires explicit auditing of performance stratified by demographic and clinical subgroups. Disaggregated evaluation should include measures of calibration, false positive rates, and failure modes, not merely aggregate accuracy.

Governance frameworks for adaptive state space networks must address data privacy, model accountability, and institutional responsibility. Federated learning offers a mechanism for training across multiple institutions without directly sharing patient data [15], and it has been applied to medical imaging collaborations to improve model generalizability while preserving privacy [16]. However, federated systems are not automatically private. Differential privacy mechanisms can provide formal guarantees, but they may interact with state space architectures in ways that require careful tuning [17]. Cross-view fusion in federated settings may also raise new risks because updates from participating sites can leak information about local data distributions through model parameters. Governance must therefore include threat modeling, secure aggregation protocols, and ongoing privacy audits. Regulatory bodies must consider whether existing medical device frameworks are adequate for models whose state transitions adapt dynamically to each patient and each view sequence. Continuous learning systems that update after deployment require additional oversight because their behavior may change over time in ways that are difficult to anticipate.

Accountability in multi-view diagnosis requires that model outputs be traceable to specific inputs and processing steps. Adaptive state space networks should be instrumented to record version identifiers, input metadata, view availability, preprocessing parameters, and fusion decisions where feasible. This information can support retrospective review when clinical outcomes are inconsistent with model predictions. Institutional review boards, clinical engineering departments, and data governance committees need clear documentation of model limitations and intended use. Algorithmic impact assessments can help identify potential harms before deployment, and post-market surveillance systems can monitor performance drift. Governance is not solely a regulatory burden; it is a design requirement that shapes data collection, model architecture, and interface design. A model that cannot be audited or explained will face substantial barriers to clinical adoption regardless of its predictive performance.

## 7. Sustainability and Lifecycle Considerations

The sustainability of adaptive state space networks depends on their computational requirements, energy consumption, and long-term maintenance. Deep learning models, particularly large transformer architectures, have raised concerns about carbon emissions and environmental impact [20]. State space models may offer efficiency advantages because they avoid quadratic attention complexity, but this does not guarantee sustainability across the entire lifecycle. Training multi-view models on large medical imaging datasets still requires substantial compute, especially when hyperparameter search, ensemble construction, and repeated cross-validation are included. The choice of patch size, state dimension, number of layers, and fusion granularity influences energy use. Model compression through pruning, quantization, and knowledge distillation can reduce inference costs, but compression may also affect robustness and fairness if not evaluated carefully. Sustainability must be considered alongside diagnostic quality rather than treated as an afterthought.

Lifecycle considerations include data versioning, model monitoring, and decommissioning. Medical imaging protocols evolve as equipment is replaced and clinical guidelines change. A model trained on older scanners may experience performance degradation when applied to newer devices. Continuous monitoring of input distributions and output confidence can detect drift, but monitoring itself consumes computational and human resources. Adaptive state space networks may be updated periodically with new data, but routine retraining raises questions about regulatory approval and clinical continuity. Version control systems for models and datasets are essential for reproducibility and for understanding how changes in training data affect downstream behavior. Model retirement is often overlooked but is necessary when performance falls below acceptable thresholds or when a replacement system is introduced. Sustainable deployment therefore requires a plan for the full operational lifecycle, not just initial training and validation.

Resource disparities across healthcare institutions also shape sustainability. Large academic medical centers may have dedicated artificial intelligence teams, high-performance computing clusters, and data engineering staff, while smaller community hospitals may lack these resources. A state space model that requires specialized hardware or complex preprocessing pipelines may widen existing inequities in diagnostic quality. Cloud-based inference can reduce local infrastructure burdens but introduces dependency on external vendors and network reliability. Edge deployment may improve latency and privacy but may require compressed models with reduced capacity. These trade-offs must be evaluated in light of the populations served. A sustainable system is one that can be maintained equitably across different institutional contexts, including those with constrained budgets and limited technical personnel.

## 8. Future Directions and Open Challenges

Several open challenges must be addressed before adaptive state space networks can achieve widespread clinical impact. First, there is a need for standardized benchmarks that evaluate multi-view robustness under realistic missing data, domain shift, and annotation noise. Current benchmark suites often assume complete views and curated datasets, which do not reflect the messy conditions of clinical practice. Benchmarking should include multi-institutional data, varied acquisition protocols, and deliberately induced view dropout. Evaluation metrics should capture calibration, fairness, interpretability, and resource consumption, not only accuracy. Such benchmarks would enable more meaningful comparisons across architectures and fusion strategies.

Second, multimodal integration beyond images should be pursued. Adaptive state space models can be extended to incorporate clinical text, laboratory values, genomic data, and longitudinal records. The challenge is to develop unified representations that respect the privacy and regulatory constraints associated with each data type. Causal representation learning may offer a path toward disentangling spurious correlations from genuine diagnostic relationships, but causal inference in high-dimensional imaging remains difficult. State space models with selective updating may be particularly well suited to learning temporal dynamics across patient histories, yet this potential remains underexplored.

Third, human-computer collaboration must be studied more carefully. Clinicians do not simply accept model outputs; they integrate them with their own expertise, patient preferences, and institutional protocols. Future systems should be designed to support collaborative reasoning, for example by highlighting conflicting views or indicating when additional imaging might reduce uncertainty. Adaptive state space networks could provide dynamic confidence estimates that guide clinician attention. However, such systems require careful interface design and empirical validation in clinical workflows. Randomized trials that measure diagnostic outcomes, clinician workload, and patient safety are necessary, but they are expensive and methodologically complex.

Fourth, governance mechanisms must evolve alongside technical advances. Regulatory frameworks for adaptive algorithms, continuous learning systems, and federated deployments remain underdeveloped. Standards bodies should work with researchers, clinicians, and patient advocates to define transparency requirements, audit procedures, and liability allocation. The connection between technical properties such as selective state updates and broader societal outcomes must be made explicit. Multi-view medical imaging models are embedded in healthcare systems that are already marked by unequal access and resource constraints. Research should aim not only to improve average performance but also to reduce disparities and avoid exacerbating existing harms. This requires interdisciplinary collaboration across computer science, medicine, ethics, law, and health policy.

## 9. Conclusion

Adaptive state space networks offer a promising architectural direction for robust multi-view medical image analysis. Their selective state transitions can integrate heterogeneous imaging evidence while avoiding the computational burden of dense attention. However, technical performance alone is insufficient for clinical translation. This paper has argued that the design and deployment of such networks must be understood as a sociotechnical problem involving infrastructure, governance, robustness, fairness, interpretability, and sustainability. Structural trade-offs between early, late, and adaptive fusion have implications for missing data handling, auditability, and institutional maintenance. The requirement for uncertainty estimation and calibration is as important as predictive accuracy in clinical settings where decisions carry high stakes. Fairness and governance frameworks must address demographic bias, privacy risks, and accountability for adaptive behavior. Sustainability demands attention to energy use, lifecycle management, and equitable access to computational resources. Future research should pursue realistic benchmarking, multimodal integration, causal reasoning, and stronger integration with clinical workflows. Only through such systemic thinking can adaptive state space networks fulfill their potential to improve diagnostic care while respecting the complexity of healthcare organizations and the needs of patients.

## References

1. Gu, A., & Dao, T. (2023). Mamba: Linear-time sequence modeling with selective state spaces. arXiv preprint arXiv:2312.00752.

2. Zhu, L., Liao, B., Zhang, Q., Wang, X., Liu, W., & Wang, X. (2024). Vision mamba: Efficient visual representation learning with bidirectional state space model. arXiv preprint arXiv:2401.09417.

3. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In International Conference on Medical Image Computing and Computer-Assisted Intervention (pp. 234–241). Springer.

4. Isensee, F., Jaeger, P. F., Kohl, S. A. A., Petersen, J., & Maier-Hein, K. H. (2021). nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. Nature Methods, 18(2), 203–211.

5. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In Advances in Neural Information Processing Systems (pp. 5998–6008).

6. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., & Houlsby, N. (2021). An image is worth 16x16 words: Transformers for image recognition at scale. In International Conference on Learning Representations.

7. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., & Guo, B. (2021). Swin Transformer: Hierarchical vision transformer using shifted windows. In Proceedings of the IEEE/CVF International Conference on Computer Vision (pp. 10012–10022).

8. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 770–778).

9. Huang, G., Liu, Z., van der Maaten, L., & Weinberger, K. Q. (2017). Densely connected convolutional networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 4700–4708).

10. Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. Nature, 542(7639), 115–118.

11. Rajpurkar, P., Irvin, J., Zhu, K., Yang, B., Mehta, H., Duan, T., Ding, D., Bagul, A., Langlotz, C., Shpanskaya, K., Lungren, M. P., & Ng, A. Y. (2017). CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning. arXiv preprint arXiv:1711.05225.

12. 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 (pp. 618–626).

13. Zheng, X., Chen, X., Gong, S., Griffin, X., & Slabaugh, G. (2025, September). Xfmamba: Cross-fusion mamba for multi-view medical image classification. In International Conference on Medical Image Computing and Computer-Assisted Intervention (pp. 672-682). Cham: Springer Nature Switzerland.

14. Kendall, A., & Gal, Y. (2017). What uncertainties do we need in Bayesian deep learning for computer vision? In Advances in Neural Information Processing Systems (pp. 5574–5584).

15. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (pp. 1273–1282). PMLR.

16. Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., Milletari, F., Horowitz, S. W., Fischer, G. S., Barash, E., Fink, G., Chaudhari, P., et al. (2020). Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. Scientific Reports, 10(1), 12598.

17. 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 (pp. 308–318).

18. 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.

19. Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. California Law Review, 104(3), 671–732.

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 (pp. 3645–3650).