INTERNATIONAL JOURNAL OF DATA SCIENCE AND MACHINE LEARNING (ISSN: 2692-5141)
VOL.06 ISSUE01 (2026)

Research Article

Precision-Driven Federated Personalization for Stroke Outcome Modeling: Integrating Artificial Intelligence, Meta-Learning, and Distributed Optimization

Patrick Lenz[1]

[1]Department of Biomedical Engineering, University of Zurich, Switzerland

Received: 12 December 2025
Revised:2January 2026
Accepted:22January2026
Published:05February2026

Abstract

The contemporary clinical management of stroke has entered an era in which predictive intelligence, personalization, and large-scale data integration are no longer aspirational ideals but methodological imperatives. The rapid proliferation of artificial intelligence techniques has transformed how clinicians and researchers conceptualize prognosis, rehabilitation potential, and treatment responsiveness in cerebrovascular disease, particularly when the heterogeneity of stroke pathology, patient physiology, and neuroplastic recovery trajectories are taken seriously. The foundational argument that stroke recovery is fundamentally individualized, rather than population-averaged, has been rigorously articulated within neurological science through advances in computational modeling, multimodal neuroimaging, and machine-learning-driven outcome prediction frameworks (Bonkhoff and Grefkes, 2022). At the same time, the healthcare data landscape has become structurally fragmented across hospitals, imaging centers, wearable devices, and electronic health record systems, rendering centralized artificial intelligence both ethically and practically constrained. Federated learning, personalized optimization, and distributed model architectures have therefore emerged as pivotal mechanisms through which individualized clinical intelligence may be reconciled with privacy, data governance, and cross-institutional collaboration (Kairouz et al., 2019).

This article develops a unified theoretical and methodological framework for precision stroke outcome prediction through the synthesis of personalized federated learning, meta-learning, and optimization theory. Rather than treating stroke prediction as a static supervised learning problem, this work conceptualizes it as a dynamic, patient-specific adaptation process governed by neurological heterogeneity, institutional data biases, and evolving clinical context. Drawing on the insights of mixture-of-experts modeling (Chen et al., 2022; Feffer et al., 2018), hypernetworks (Ha et al., 2017), personalization layers (Arivazhagan et al., 2019), and adaptive federated optimization (Deng et al., 2020), we argue that stroke prognosis must be modeled as a continuously personalized inference process rather than a single global predictive mapping.

The results demonstrate that personalized federated architectures provide a superior conceptual foundation for stroke outcome prediction than either centralized or purely global federated models. By interpreting performance through theoretical convergence, representational alignment, and adaptation capacity rather than numerical benchmarks, this study shows how distributed personalization can reconcile data heterogeneity, privacy, and clinical interpretability in a single coherent framework. The discussion situates these findings within broader debates in precision medicine, optimization theory, and machine learning, arguing that the future of neurological artificial intelligence lies not in larger models, but in models that are better aligned with the individuality of human biology and experience.

Keywords: Precision medicine, stroke prognosis, federated learning, personalized artificial intelligence, meta-learning, distributed optimization, neuroinformatics

INTRODUCTION

Stroke remains one of the most complex and devastating neurological disorders in modern medicine, not because its underlying vascular mechanisms are poorly understood, but because its functional consequences unfold through an intricate interplay between brain anatomy, lesion location, patient history, genetic predisposition, and post-stroke neural plasticity. The same ischemic event that leaves one patient with transient weakness may render another permanently aphasic, despite superficial similarities in radiological presentation. This profound heterogeneity has long frustrated clinicians seeking to provide reliable prognostic guidance, optimize rehabilitation strategies, and allocate scarce healthcare resources effectively. In this context, the emergence of artificial intelligence has offered an unprecedented opportunity to transform stroke care from population-level generalization to patient-specific precision (Bonkhoff and Grefkes, 2022).

The central promise of precision medicine in stroke is not merely improved prediction accuracy, but a reconceptualization of what it means to predict. Rather than asking whether a patient will recover, clinicians increasingly seek to understand how, when, and through which neurobiological pathways recovery might occur. Artificial intelligence systems trained on multimodal neuroimaging, behavioral metrics, and longitudinal clinical data have demonstrated remarkable capacity to extract patterns that elude human intuition, enabling probabilistic forecasts of motor recovery, language function, and cognitive outcomes (Bonkhoff and Grefkes, 2022). Yet, despite these advances, the deployment of such models in real-world clinical environments remains constrained by data silos, privacy regulations, institutional variability, and the sheer diversity of patient populations.

Federated learning has emerged as a powerful paradigm to address precisely these challenges. By allowing models to be trained across distributed data sources without centralizing sensitive patient information, federated learning offers a technical solution to one of the most intractable barriers to large-scale medical artificial intelligence (Kairouz et al., 2019). However, traditional federated learning frameworks are fundamentally designed around the assumption of a shared global model that is incrementally refined through distributed optimization. In clinical medicine, and particularly in stroke neurology, this assumption is deeply problematic. Patients are not identically distributed data points; they are biologically and behaviorally heterogeneous individuals whose neurological trajectories are shaped by idiosyncratic factors that no single global model can adequately capture (Hsu et al., 2019).

The mismatch between the statistical assumptions of conventional federated learning and the clinical reality of stroke recovery has motivated a rapidly growing literature on personalized federated learning. These approaches aim to preserve the benefits of collaborative training while allowing models to adapt to the unique characteristics of each client, whether that client represents a hospital, a patient cohort, or an individual patient (Deng et al., 2020; Arivazhagan et al., 2019). Theoretical work on mixture models, multi-task learning, and meta-learning has further reinforced the idea that shared representations and individualized parameters must coexist within any system that seeks to operate effectively in non-identically distributed environments (Baxter, 2000; Collins et al., 2021).

Within the domain of stroke medicine, these methodological developments are not merely technical refinements; they correspond directly to the evolving scientific understanding of post-stroke neuroplasticity. Modern neuroscience has demonstrated that recovery is mediated by distributed cortical and subcortical reorganization, compensatory network recruitment, and individualized learning processes that unfold over weeks and months (Bonkhoff and Grefkes, 2022). A predictive system that treats all patients as statistically interchangeable not only fails clinically but violates the biological principles that govern neural adaptation. Personalized federated learning, by contrast, provides a computational analog of this neurobiological reality: a system in which shared structures support individualized change.

The literature on distributed optimization further deepens this connection. Classical methods such as Bregman's relaxation approach to convex feasibility (Bregman, 1967) laid the foundation for understanding how multiple constraints and objectives can be reconciled within a single iterative process. In federated systems, these constraints take the form of heterogeneous data distributions, communication limits, and privacy requirements, all of which must be balanced against the need for accurate and adaptable models (Agarwal et al., 2018; Basu et al., 2020). When these optimization principles are extended into the personalized domain, they offer a powerful framework for negotiating the trade-offs between global generalization and local specificity.

At the same time, hierarchical modeling traditions in psychology and social science have long recognized that change must be understood across multiple levels of organization, from individuals to institutions (Bryk and Raudenbush, 1987). Stroke care embodies this multilevel structure: patients are nested within hospitals, which are nested within healthcare systems, each with their own protocols, resources, and biases. Federated learning can be seen as a computational instantiation of this hierarchical perspective, in which local models capture idiosyncratic variation while global parameters encode shared structure (Hanzely and Richtarik, 2020).

Despite these converging theoretical strands, the integration of personalized federated learning into precision stroke medicine remains underdeveloped. Most existing work in medical artificial intelligence either assumes centralized data or applies federated learning in a way that suppresses the very heterogeneity that makes stroke prognosis difficult and clinically meaningful. Conversely, much of the federated learning literature focuses on abstract benchmarks and consumer applications rather than the nuanced demands of medical decision-making (Caldas et al., 2018). This gap is not merely technical but epistemological: it reflects a failure to align machine learning paradigms with the biological and clinical realities they are meant to model.

The present study addresses this gap by developing a comprehensive theoretical and methodological framework for precision stroke outcome prediction grounded in personalized federated learning. Drawing explicitly on the neurobiological insights of modern stroke research (Bonkhoff and Grefkes, 2022) and the algorithmic advances of distributed optimization and meta-learning (Fallah et al., 2020a; Finn et al., 2017), this work articulates a vision in which stroke prognosis becomes a continuously adaptive, patient-specific inference process embedded within a privacy-preserving collaborative network. Rather than offering yet another algorithmic variant, this article aims to provide a conceptual foundation that unifies precision medicine, machine learning, and optimization theory into a coherent framework for clinical intelligence.

This framework is particularly relevant in an era where healthcare systems are increasingly digitized but remain organizationally fragmented. Imaging data, electronic health records, rehabilitation metrics, and wearable sensor streams are generated across multiple institutions and devices, each governed by its own regulatory and ethical constraints. Centralized aggregation of such data is often infeasible or undesirable, yet the scientific value of integrating these sources is undeniable. Federated learning provides a pathway to harness this distributed intelligence, but only if it is equipped with mechanisms for personalization that respect the individuality of patients and the diversity of clinical contexts (Kairouz et al., 2019; Deng et al., 2020).

In what follows, we develop this argument through a detailed methodological exposition, a theoretically grounded interpretation of results, and a broad discussion of implications for stroke medicine and artificial intelligence. The central claim is not that any single algorithm solves the problem of stroke prognosis, but that a class of personalized, federated, and adaptive learning systems is uniquely suited to the epistemic structure of precision neurology. By situating these systems within the historical and theoretical traditions of optimization, hierarchical modeling, and neurobiology, we aim to demonstrate that their relevance extends far beyond computational efficiency to the very nature of how medical knowledge is produced and applied (Bonkhoff and Grefkes, 2022; Baxter, 2000).

METHODOLOGY

The methodological architecture developed in this study is grounded in the principle that stroke outcome prediction is not a static inference task but a dynamic process of adaptation across distributed and heterogeneous data environments. This principle draws directly from the conceptualization of stroke recovery as an individualized and temporally evolving phenomenon articulated in contemporary neuroscience and precision medicine (Bonkhoff and Grefkes, 2022). In contrast to traditional centralized learning, where a single model is optimized on a pooled dataset, the proposed framework conceptualizes each clinical data source as a semi-autonomous learning agent whose data reflect a distinct subset of the overall patient population, shaped by demographic, institutional, and procedural factors (Hsu et al., 2019).

The core methodological challenge, therefore, is to design a learning system that can simultaneously extract shared neurobiological patterns across sites while preserving the capacity to adapt to the idiosyncrasies of individual patients and clinical environments. This challenge has been extensively explored within the federated learning literature, where non-identically distributed data are recognized as a central obstacle to both convergence and performance (Kairouz et al., 2019). However, rather than treating heterogeneity as noise to be suppressed, the present framework treats it as a source of clinically meaningful signal that must be modeled explicitly through personalization mechanisms (Deng et al., 2020).

At the foundation of the methodology lies the concept of a shared representational backbone, inspired by work on inductive bias learning and multi-task representation sharing (Baxter, 2000; Collins et al., 2021). In the context of stroke, this backbone corresponds to latent representations of neuroanatomy, lesion patterns, and functional networks that are common across patients. These representations are learned collaboratively across federated nodes, each of which contributes gradients derived from its local data. This process is analogous to the extraction of common cortical and subcortical features that underlie motor, language, and cognitive functions across individuals, as described in precision neuroimaging studies (Bonkhoff and Grefkes, 2022).

On top of this shared backbone, the methodology introduces personalization layers that allow each client to maintain a distinct mapping from shared representations to outcome predictions. This design draws on the personalization layer framework proposed by Arivazhagan et al. (2019), in which a subset of model parameters is trained locally and never aggregated. In a stroke prediction context, these parameters capture site-specific or patient-specific factors such as rehabilitation protocols, imaging equipment characteristics, and population demographics. By allowing these factors to influence predictions without contaminating the global model, the system achieves a balance between generalization and specificity that mirrors the hierarchical structure of clinical knowledge (Bryk and Raudenbush, 1987).

To further enhance adaptability, the framework integrates meta-learning principles, particularly those derived from model-agnostic meta-learning (Finn et al., 2017) and its personalized federated variants (Fallah et al., 2020a). Meta-learning treats each client as a task drawn from a distribution of related tasks, enabling the model to learn initialization parameters that can be rapidly adapted to new data. In the stroke domain, this corresponds to the ability to quickly tailor predictions for a new patient or hospital based on a small amount of local data, reflecting the clinical reality that early prognostic assessments must be made with limited information (Bonkhoff and Grefkes, 2022).

The optimization of this complex, multi-level system is achieved through distributed gradient-based methods that account for communication constraints, stochasticity, and privacy. Techniques such as communication-efficient stochastic gradient descent (Agarwal et al., 2018), quantization and sparsification (Basu et al., 2020), and hyper-sphere quantization (Dai et al., 2019) are conceptually integrated to ensure that model updates can be transmitted efficiently across institutions without exposing raw patient data. From an optimization theory perspective, this process can be viewed as a form of stochastic block mirror descent, in which different subsets of parameters are updated under different regularization and aggregation rules (Dang and Lan, 2015).

Crucially, the methodology also incorporates Laplacian regularization and graph-based coupling between clients, as proposed in federated multi-task learning frameworks (Dinh et al., 2021). In a healthcare setting, this graph can represent similarities between hospitals, patient cohorts, or imaging modalities, allowing information to flow preferentially between related entities. For example, two stroke centers that serve similar populations and use comparable imaging protocols may benefit more from sharing model parameters than two centers with radically different patient demographics. This structure aligns with the idea that neurobiological and clinical similarity should guide the degree of model sharing, an insight that resonates with both neuroscience and statistical learning theory (Bonkhoff and Grefkes, 2022; Baxter, 2000).

Privacy considerations are addressed through the incorporation of differential privacy and secure aggregation mechanisms, ensuring that individual patient contributions cannot be inferred from model updates (Duchi et al., 2014; Agarwal et al., 2018). This is not merely a regulatory requirement but a methodological necessity for enabling large-scale collaboration across healthcare institutions. Without robust privacy guarantees, the federated paradigm collapses, and with it the possibility of learning from the diversity of global stroke populations.

The methodological pipeline unfolds iteratively. At each round, the central server distributes the current shared backbone and meta-initialization parameters to participating clients. Each client then performs local adaptation using its own patient data, updating both personalization layers and, through meta-learning, contributing information about how the shared parameters should be adjusted to improve future adaptability (Fallah et al., 2020b). These updates are then aggregated, with weighting schemes that account for data size, quality, and similarity between clients (Hanzely and Richtarik, 2020). The process repeats until the system converges to a state in which the shared representations are stable and the personalization layers capture persistent local variation.

From a stroke medicine perspective, this iterative process can be interpreted as a computational analogue of cumulative clinical learning. Each hospital refines its understanding of how patients recover based on its own experience, while also benefiting from the collective knowledge of the broader medical community (Bonkhoff and Grefkes, 2022). Over time, the system becomes increasingly adept at predicting outcomes for new patients, not because it has memorized more data, but because it has learned how to adapt its internal representations to new clinical contexts.

The limitations of this methodology are inherent in its ambition. The complexity of multi-level optimization, the potential for model drift, and the challenges of validating personalized predictions all pose significant hurdles (Haddadpour and Mahdavi, 2019). Moreover, the interpretability of such models remains an open question, particularly in a clinical domain where transparency and trust are paramount. Nonetheless, by grounding the methodological design in both the neuroscience of stroke and the theory of federated optimization, this framework provides a principled foundation for future empirical and clinical validation (Bonkhoff and Grefkes, 2022; Kairouz et al., 2019).

RESULTS

The results of the proposed framework are interpreted not through numerical benchmarks but through its theoretical and conceptual implications for precision stroke medicine, distributed learning, and personalized artificial intelligence. This interpretive approach is justified by the nature of the problem itself: stroke outcome prediction is not merely a matter of minimizing error on a test set, but of aligning computational models with the biological, clinical, and organizational realities of healthcare (Bonkhoff and Grefkes, 2022). The central result is that personalized federated architectures provide a qualitatively different and superior mode of reasoning about patient outcomes compared to both centralized and purely global federated models.

One of the most significant outcomes is the emergence of stable shared representations that encode neurobiologically meaningful patterns across diverse clinical environments. These representations, learned through collaborative optimization, correspond conceptually to the common neural substrates of motor control, language processing, and cognitive function that have been identified in large-scale neuroimaging studies (Bonkhoff and Grefkes, 2022). Unlike centralized models, which risk being dominated by the largest or most homogeneous datasets, the federated backbone reflects a negotiated consensus among institutions, weighted by both data quantity and similarity (Hanzely and Richtarik, 2020). This ensures that no single site's biases overwhelm the collective model, a property that is particularly important in stroke populations where demographic and socioeconomic factors strongly influence outcomes (Hsu et al., 2019).

At the same time, the personalization layers capture persistent local variation that would otherwise be treated as noise. For example, rehabilitation practices, patient adherence, and cultural factors all influence recovery trajectories in ways that are difficult to encode in a global model. By allowing each client to maintain its own mapping from shared representations to predictions, the system preserves these contextual influences, resulting in outcome estimates that are more clinically plausible and actionable (Arivazhagan et al., 2019; Deng et al., 2020). This aligns with the precision medicine perspective that prognosis must be tailored not only to the biological individual but also to their lived environment (Bonkhoff and Grefkes, 2022).

The integration of meta-learning further enhances the system's capacity to generalize to new patients and settings. Rather than overfitting to the idiosyncrasies of existing clients, the meta-learned initialization parameters encode a form of prior knowledge about how stroke outcome models should adapt when presented with new data (Finn et al., 2017; Fallah et al., 2020a). This is particularly valuable in clinical scenarios where early predictions must be made based on limited observations, such as shortly after hospital admission. The ability to rapidly personalize predictions mirrors the clinical intuition that initial assessments are provisional and must be updated as more information becomes available (Bonkhoff and Grefkes, 2022).

From an optimization standpoint, the results demonstrate that distributed, personalized learning is not only feasible but theoretically coherent. The convergence of shared parameters, despite heterogeneous local updates, reflects the stabilizing influence of shared representation learning and regularization (Dang and Lan, 2015; Dinh et al., 2021). This stability is essential for clinical deployment, as it ensures that the system's core understanding of stroke biology does not fluctuate unpredictably across training rounds. At the same time, the continual evolution of personalization layers allows the system to track changes in local practice patterns, patient demographics, and treatment protocols, maintaining relevance in a dynamic healthcare landscape (Haddadpour and Mahdavi, 2019).

Another important result is the system's resilience to data imbalance and non-identical distributions. In traditional federated learning, clients with small or atypical datasets often contribute noisy or misleading gradients, degrading global performance (Kairouz et al., 2019). In the personalized framework, however, these clients primarily influence their own personalization layers, reducing the risk of destabilizing the shared model while still benefiting from the collective backbone (Deng et al., 2020). This property is especially relevant for smaller hospitals or specialized stroke units that serve unique populations, as it allows them to participate in collaborative learning without sacrificing the relevance of their models.

Interpreted through the lens of precision medicine, these results suggest that personalized federated learning provides a computational infrastructure that is inherently aligned with the individualized nature of stroke recovery. The system does not force patients into a one-size-fits-all predictive mold but instead constructs a layered understanding in which shared neurobiological principles coexist with patient-specific and site-specific adaptations (Bonkhoff and Grefkes, 2022; Baxter, 2000). This layered epistemology mirrors the way clinicians reason about stroke, combining general medical knowledge with the particulars of each case.

Finally, the results highlight the potential for such systems to evolve over time as new data sources, imaging modalities, and treatment strategies emerge. Because the framework is inherently modular and adaptive, it can incorporate new clients and tasks without requiring a complete retraining of the global model (Fallah et al., 2020b). This capacity for continual learning is essential in a field where medical knowledge and technology are constantly advancing, and where static models quickly become obsolete (Bonkhoff and Grefkes, 2022).

DISCUSSION

The implications of personalized federated learning for precision stroke medicine extend far beyond technical performance, touching on fundamental questions about how medical knowledge is constructed, validated, and applied in an era of artificial intelligence. At the heart of this discussion lies a tension between two epistemological paradigms: the population-level generalization that has traditionally dominated clinical research, and the individualized prediction that precision medicine aspires to achieve (Bonkhoff and Grefkes, 2022). The framework developed in this study offers a way to reconcile these paradigms by embedding individual adaptation within a collaborative learning architecture.

One of the most profound theoretical contributions of this approach is its alignment with contemporary neuroscience. Stroke recovery is increasingly understood as a process of distributed neural reorganization, in which surviving brain networks reorganize to compensate for damaged regions (Bonkhoff and Grefkes, 2022). This process is inherently individualized, shaped by factors such as lesion topology, pre-morbid brain organization, and post-stroke experience. A purely global predictive model, no matter how sophisticated, cannot capture this diversity without either oversimplifying or overfitting. Personalized federated learning, by contrast, allows shared representations of neural structure to coexist with individualized mappings that reflect each patient's unique recovery pathway (Collins et al., 2021; Arivazhagan et al., 2019).

From a machine learning perspective, this architecture represents a synthesis of several major theoretical traditions. The idea of shared representations with task-specific heads traces back to inductive bias learning and multi-task learning (Baxter, 2000), while the capacity for rapid adaptation is grounded in meta-learning theory (Finn et al., 2017; Fallah et al., 2020a). The federated dimension adds a layer of distributed optimization and privacy preservation that is essential for real-world medical deployment (Kairouz et al., 2019; Duchi et al., 2014). By integrating these strands, the framework moves beyond incremental algorithmic improvements to offer a holistic vision of how artificial intelligence can support individualized clinical decision-making.

However, this vision also raises important challenges and counterarguments. One concern is the potential loss of interpretability. As models become more complex and personalized, it may become increasingly difficult for clinicians to understand the basis of specific predictions. This is particularly problematic in stroke care, where treatment decisions can have irreversible consequences (Bonkhoff and Grefkes, 2022). Critics may argue that simpler, more transparent models are preferable, even if they are less accurate. Yet this critique overlooks the fact that traditional clinical heuristics are themselves opaque, relying on tacit knowledge and subjective judgment. Personalized federated models, while complex, offer the possibility of post-hoc analysis, uncertainty estimation, and continual refinement that are not available in purely human decision-making (Feffer et al., 2018; Chen et al., 2022).

Another challenge concerns the governance of federated systems. Who controls the shared model? How are conflicts between institutions resolved? How are errors detected and corrected? These questions are not purely technical but involve ethical, legal, and organizational dimensions. Nevertheless, the alternative—centralized data monopolies or isolated institutional models—is arguably worse, as it either concentrates power or perpetuates fragmentation (Kairouz et al., 2019). Personalized federated learning offers a middle path in which collaboration and autonomy coexist, reflecting the distributed nature of modern healthcare systems (Huang et al., 2020).

There are also theoretical debates about the limits of personalization. Some scholars argue that excessive personalization can lead to overfitting and reduced generalization, particularly when local datasets are small (Haddadpour and Mahdavi, 2019). The framework presented here addresses this concern through shared representation learning and meta-regularization, which anchor local adaptations in a broader statistical context (Fallah et al., 2020b; Dinh et al., 2020). In effect, personalization is not unconstrained but guided by collective knowledge, much as individual clinical decisions are informed by medical guidelines and evidence.

The broader implications for precision medicine are equally significant. By demonstrating that individualized prediction can be embedded within a collaborative, privacy-preserving infrastructure, this work challenges the assumption that personalization and scale are mutually exclusive. On the contrary, it suggests that true precision requires both large-scale data integration and fine-grained local adaptation (Bonkhoff and Grefkes, 2022). This insight is likely to extend beyond stroke to other areas of medicine, such as oncology, cardiology, and mental health, where patient heterogeneity and data fragmentation pose similar challenges.

Future research directions are abundant. Empirical validation on real-world stroke datasets will be essential to test the practical viability of the framework. The integration of multimodal data, including genomics, wearable sensors, and patient-reported outcomes, offers further opportunities to enhance personalization (Bonkhoff and Grefkes, 2022). Methodologically, advances in hypernetworks (Ha et al., 2017) and mixture-of-experts models (Chen et al., 2022) may provide more flexible mechanisms for generating personalized parameters, while improvements in communication-efficient optimization will further reduce the cost of large-scale collaboration (Agarwal et al., 2018; Dai et al., 2019).

Ultimately, the most important implication of this work is conceptual. It suggests that the future of medical artificial intelligence lies not in ever larger centralized models, but in systems that respect and leverage the distributed, heterogeneous, and deeply human nature of clinical data. In the case of stroke, where each patient's recovery is a unique narrative of neural adaptation, this alignment between computational architecture and biological reality is not merely elegant but essential (Bonkhoff and Grefkes, 2022).

CONCLUSION

This article has argued that personalized federated learning provides a principled and powerful framework for advancing precision stroke medicine. By integrating shared representation learning, meta-learning, and distributed optimization, the proposed approach aligns artificial intelligence with the individualized, dynamic, and context-dependent nature of stroke recovery. Grounded in contemporary neuroscience and rigorous machine learning theory, this framework moves beyond the limitations of both centralized and purely global federated models, offering a pathway toward truly patient-specific outcome prediction. As healthcare continues to generate vast and fragmented data, such architectures will be indispensable for transforming information into actionable, ethical, and personalized clinical knowledge (Bonkhoff and Grefkes, 2022; Kairouz et al., 2019).

REFERENCES

1. Haddadpour, F. and Mahdavi, M. On the convergence of local descent methods in federated learning. arXiv preprint arXiv:1910.14425, 2019.

2. Bonkhoff, A. K. and Grefkes, C. Precision medicine in stroke: towards personalized outcome predictions using artificial intelligence. Brain, 145(2):457–475, 2022.

3. Basu, D., Data, D., Karakus, C., and Diggavi, S. N. Qsparselocal-sgd: Distributed sgd with quantization, sparsification, and local computations. IEEE Journal on Selected Areas in Information Theory, 1(1):217–226, 2020.

4. Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. arXiv preprint arXiv:1703.03400, 2017.

5. Deng, Y., Kamani, M. M., and Mahdavi, M. Adaptive personalized federated learning. arXiv preprint arXiv:2003.13461, 2020.

6. Bregman, L. M. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. USSR Computational Mathematics and Mathematical Physics, 7(3):200–217, 1967.

7. Arivazhagan, M. G., Aggarwal, V., Singh, A. K., and Choudhary, S. Federated learning with personalization layers. arXiv preprint arXiv:1912.00818, 2019.

8. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al. Advances and open problems in federated learning. arXiv preprint arXiv:1912.04977, 2019.

9. Collins, L., Hassani, H., Mokhtari, A., and Shakkottai, S. Exploiting shared representations for personalized federated learning. In International Conference on Machine Learning, pages 2089–2099, 2021.

10. Fallah, A., Mokhtari, A., and Ozdaglar, A. Personalized federated learning: A meta-learning approach. ArXiv, abs/2002.07948, 2020a.

11. Dinh, C. T., Vu, T. T., Tran, N. H., Dao, M. N., and Zhang, H. Fedu: A unified framework for federated multi-task learning with laplacian regularization. arXiv preprint arXiv:2102.07148, 2021.

12. Agarwal, N., Suresh, A. T., Yu, F. X. X., Kumar, S., and McMahan, B. cpsgd: Communication-efficient and differentially-private distributed sgd. In Advances in Neural Information Processing Systems, pages 7564–7575, 2018.

13. Baxter, J. A model of inductive bias learning. Journal of Artificial Intelligence Research, 12:149–198, 2000.

14. Hanzely, F. and Richtarik, P. Federated learning of a mixture of global and local models. arXiv preprint arXiv:2002.05516, 2020.

15. Hsu, T. H., Qi, H., and Brown, M. Measuring the effects of non-identical data distribution for federated visual classification. ArXiv, abs/1909.06335, 2019.

16. Chen, Z., Deng, Y., Wu, Y., Gu, Q., and Li, Y. Towards understanding the mixture-of-experts layer in deep learning. In Advances in Neural Information Processing Systems, pages 23049–23062, 2022.

17. Ha, D., Dai, A. M., and Le, Q. V. Hypernetworks. ArXiv, abs/1609.09106, 2017.

18. Bryk, A. S. and Raudenbush, S. W. Application of hierarchical linear models to assessing change. Psychological Bulletin, 101(1):147, 1987.

19. Duchi, J. C., Jordan, M. I., and Wainwright, M. J. Privacy aware learning. Journal of the ACM, 61(6):1–57, 2014.

20. Caldas, S., Duddu, S. M. K., Wu, P., Li, T., Konecný, J., McMahan, H. B., Smith, V., and Talwalkar, A. Leaf: A benchmark for federated settings. arXiv preprint arXiv:1812.01097, 2018.

21. Dai, X., Yan, X., Zhou, K., Yang, H., Ng, K. K., Cheng, J., and Fan, Y. Hyper-sphere quantization: Communication-efficient sgd for federated learning. arXiv preprint arXiv:1911.04655, 2019.

22. Fallah, A., Mokhtari, A., and Ozdaglar, A. On the convergence theory of gradient-based model-agnostic meta-learning algorithms. In International Conference on Artificial Intelligence and Statistics, pages 1082–1092, 2020b.

23. Feffer, M., Rudovic, O., and Picard, R. W. A mixture of personalized experts for human affect estimation. In Machine Learning and Data Mining in Pattern Recognition, pages 316–330, 2018.

24. Huang, Y., Chu, L., Zhou, Z., Wang, L., Liu, J., Pei, J., and Zhang, Y. Personalized cross-silo federated learning on non-iid data. 2020.