Advanced Artificial Intelligence Paradigms in Medical Diagnostics: Algorithmic Continual Learning, Federated Architectures, and Clinical Regulatory Frameworks

The integration of artificial intelligence into biomedical science and clinical healthcare has reached a critical technological and regulatory inflection point. Historically, the deployment of machine learning in medical diagnostics has relied on static, "locked" algorithms. These systems are trained on curated, retrospective datasets, validated, and subsequently frozen before clinical deployment. While this paradigm simplifies regulatory oversight and initial validation, it is fundamentally misaligned with the dynamic, non-stationary realities of clinical medicine. Patient demographics evolve, epidemiological landscapes shift, medical imaging hardware undergoes continuous technological refinement, and institutional treatment protocols are frequently updated. A static neural network deployed in such an environment inevitably suffers from data drift, leading to silent performance degradation and potential patient harm[1].

The computational imperative is therefore to transition toward dynamic, adaptive systems capable of lifelong, continuous learning. However, programming deep neural networks for continuous adaptation introduces a severe algorithmic vulnerability known as "catastrophic forgetting"[4]. As a neural network acquires novel information from sequential data streams, the gradient updates optimized for the new data frequently overwrite the synaptic weight representations of previously learned clinical tasks, destroying the network's foundational capabilities[4]. Resolving this stability-plasticity dilemma requires a sophisticated interdisciplinary convergence. It demands advanced neural network architectures that can dynamically expand their capacity or route information sparsely, federated learning protocols that can securely aggregate knowledge across global hospital networks without violating patient privacy, and radically modernized regulatory frameworks that shift oversight from static product clearance to continuous, algorithm-lifecycle governance[8].

The ensuing analysis provides an exhaustive examination of the modern clinical artificial intelligence ecosystem. It explores the algorithmic foundations of continual learning and catastrophic forgetting, conducts a deep technical deconstruction of dynamic neural networks through the lens of the experimental MORPH continuous-learning architecture, and evaluates the management of multimodal clinical data via Mixture of Experts (MoE) topologies[12]. Furthermore, the analysis investigates Federated Continual Learning (FCL) across decentralized healthcare networks and the hardware constraints of deploying these models at the clinical

edge[9]. Crucially, the report dissects the regulatory frameworks—specifically the United States Food and Drug Administration (FDA) Predetermined Change Control Plan (PCCP) and the Australian Therapeutic Goods Administration (TGA) Essential Principles—that render the clinical deployment of adaptive models legally and ethically viable[17]. Finally, the analysis demonstrates how localized, world-leading research ecosystems, specifically the Melbourne Biomedical AI precinct, are actively translating these advanced computational theories into precision oncology clinical trials and transformative diagnostic networks[19].

The Algorithmic Foundation: Continual Learning and the Mitigation of Catastrophic Forgetting

Continual Learning (CL), alternatively referred to as lifelong or incremental learning, is a machine learning paradigm designed to sequentially acquire new skills or adapt to new data distributions without losing previously acquired knowledge[4]. In traditional deep learning, overcoming catastrophic forgetting is typically achieved by aggregating all historical and new data and retraining the model from scratch in randomized batches[7]. In clinical realities, this is computationally prohibitive and legally complex. Data arrives in temporal streams, and historical datasets from prior years or different institutions may be archived, siloed, or restricted by evolving privacy legislation such as the Health Insurance Portability and Accountability Act (HIPAA) in the United States or the General Data Protection Regulation (GDPR) in the European Union[7].

Modalities of Clinical Data Shift

The necessity for continuous adaptation is driven by three primary categories of distribution shifts encountered in healthcare environments:

1. **Domain Shift (Spatial and Technical Drift):** This occurs when algorithms are deployed across different hospitals utilizing diverse imaging hardware or laboratory assay techniques. For example, variability in Computed Tomography (CT) scanner reconstruction kernels significantly impacts radiomics features[5]. A deep learning model trained exclusively on imaging data from Scanner A will frequently fail to generalize when processing images from Scanner B, leading to elevated false-negative rates in tasks such as lung nodule detection[24].
2. **Concept Drift (Temporal and Epidemiological Drift):** This involves the evolution of biological diseases and patient populations over time. The emergence of novel viral strains, such as the rapid evolution of COVID-19 variants, alters both the radiological presentation on chest X-rays and the predictive value of clinical biomarkers[2].
3. **Institutional and Protocol Drift:** This involves changes in clinical guidelines and hospital practices. In Intensive Care Units (ICUs), clinical decision support systems (CDSS) often suffer from drift not because underlying human physiology changes, but because treatment protocols—such as sedation practices following the adoption of the ABCDEF bundle, or evolving vasopressor recommendations—are updated based on new evidence[26]. Retraining a monolithic model on these new protocols often obscures

whether the model's updated predictions are driven by patient biology or merely mimicking the new institutional coding practices[26].

A highly illustrative demonstration of these shifts occurred during the COVID-19 pandemic. Researchers analyzing electronic health records for 143,049 patients across seven hospitals in Toronto utilized a label-agnostic monitoring pipeline employing maximum mean discrepancy testing to detect massive data shifts[2]. These shifts were driven by changes in patient demographics, admission sources (e.g., increased transfers from nursing homes), and critical laboratory assays[2]. By deploying drift-triggered continual learning strategies rather than relying on a static model, the researchers proactively updated their clinical AI, resulting in a statistically significant improvement in the Area Under the Receiver Operating Characteristic curve (AUROC) for predicting in-hospital mortality (Delta AUROC = 0.44; P = .007)[2].

Algorithmic Paradigms for CL Mitigation

To combat catastrophic forgetting while adapting to these clinical shifts, biomedical computer science has developed three primary algorithmic paradigms, frequently benchmarked using standardized datasets like MedMNIST—a diverse collection of lightweight 2D and 3D biomedical image datasets designed to evaluate stability-plasticity trade-offs[27]:

1. **Regularization-based Methods:** These algorithms, which include Elastic Weight Consolidation (EWC) and Learning without Forgetting (LwF), operate by computing the Fisher information matrix to identify which synaptic weights within the neural network are most critical to previously learned clinical tasks[4]. During the training of a new task, a mathematical penalty is applied to the loss function to constrain changes made to these highly critical parameters[7]. While regularization is highly memory-efficient because it does not require storing raw historical patient data, empirical evidence indicates it is often insufficient for severe distribution shifts[32]. In class-incremental histopathology settings, regularization methods have exhibited severe forgetting (75-100% loss of prior knowledge) due to class imbalance and logit bias toward newly introduced diagnostic categories[32].
2. **Replay and Rehearsal-based Methods:** These algorithms maintain a dynamic memory buffer containing a highly curated subset of historical data (exemplars) or utilize generative adversarial networks (GANs) to synthesize historical features during new training phases[1]. Dynamic Memory (DM) rehearsal is highly effective at anchoring the network's representations. In multi-scanner lung nodule detection studies, static sequential training yielded a negative backward transfer (BWT of -0.011), while dynamic memory approaches successfully integrated new scanners without degrading performance on legacy hardware domains[24]. The primary limitation of rehearsal in healthcare is the privacy, security, and memory constraints associated with storing raw patient imaging or genomic data[33].
3. **Architecture and Prompt-based Methods:** Architecture-based methods physically alter the network topology by allocating new, isolated neural branches for new clinical tasks (e.g., Progressive Neural Networks), preventing interference by design but resulting

in linear parameter growth over time[30]. Prompt-based learning (e.g., DualPrompt, CODA-Prompt) represents a more recent innovation adapted from natural language processing. These methods utilize a large, frozen pretrained foundation model (such as a Vision Transformer or ViT) and append learnable, task-specific "prompt" tokens to the input layer[33]. In benchmarking on the NCT-CRC-HE-100K histopathology dataset, DualPrompt demonstrated exceptional performance with minimal training epochs, serving as a highly effective, exemplar-free alternative that avoids the privacy concerns of storing raw patient data[33].

Dynamic Neural Networks and the MORPH Framework: A Technical Deconstruction

While theoretical models of continual learning provide the mathematical basis for adaptation, translating these algorithms into deployable software requires rigorous, system-level engineering. The MORPH (Continual Polymorphic Numeric Exploration) framework serves as an advanced experimental testbed that simulates the strict deployment constraints of continual learning systems in medical environments[12]. By analyzing the successive iterations of the MORPH codebase (v1, v2, and v3), we can deconstruct the exact mechanisms by which dynamic neural networks manage capacity, route inferences, and validate adaptations.

MORPH v1: Context Gating and Shadow Training

The foundational challenge addressed by MORPH v1 is the protection of an incumbent base model while safely accommodating new, context-specific knowledge. The architecture eschews monolithic retraining in favor of an additive context gate mechanism. In the Go-based implementation of MORPH v1, the network consists of a base network and an optional specialist network. During inference, the output prediction y is calculated by first passing the input through the base network, and if a specialist exists and the current clinical context matches its target, adding the specialist's output scaled by an Alpha gate parameter[12].

The controller continuously observes data streams and calculates Exponential Moving Average (EWMA) losses for distinct clinical contexts[12]. A healthy context, where the model performs below a defined error threshold (e.g., HealthyLoss: 0.02), does not trigger any growth[12]. However, if the controller detects a persistent residual error in a specific context that fails to improve over several observation windows, it initiates a "shadow training" phase[12].

During shadow training, a candidate specialist network is initialized and trained on a bounded replay window of observed data (e.g., a maximum of 512 historical samples), specifically targeting the residual error left by the frozen base model[12]. Crucially, the MORPH v1 validation protocol tests the candidate across predefined alpha gates (0.25, 0.5, 1.0) on a sequestered validation set[12]. To pass the validation gate, the candidate must achieve a minimum relative product gain (e.g., 15% improvement in the target context) while simultaneously passing a strict retention check for the old context (e.g., a maximum absolute regression of 0.02)[12].

Empirical measurements from MORPH v1 demonstrated the efficacy of this mechanism. Across

ten paired execution seeds operating on a 1,250-sample stream (Phases A, B, and C), the D_MORPH model successfully triggered protected specialist growth and passed all retention gates in 10/10 runs[12]. The dynamic framework achieved a mean product MSE of 0.97313 utilizing an average of only 24.12 parameters, whereas a pre-allocated static expert network (C_STATIC_2EXPERT) utilized 36.00 parameters throughout the entire execution[12].

Model Variant	Sum MSE	Product MSE	Mean Resident Params	Train MS
A_LINEAR	0.18456	0.69425	4.00	2.099
B_MLP_FIXED	0.032129	0.061911	40.00	5.866
C_STATIC_2EXPERT	6.973e-08	0.04449	36.00	3.059
D_MORPH	6.973e-08	0.97313	24.12	8.957

Table 1: Aggregated empirical outcomes for MORPH v1 across 10 paired seeds, highlighting parameter efficiency and error retention[12].

MORPH v2: Structural Search and Capacity Matching

While MORPH v1 utilized fixed-width specialists, MORPH v2 introduces continuous structural exploration, recognizing that the optimal capacity required for a new clinical task is rarely known a priori[13]. The v2 architecture expands the shadow training mechanism to evaluate different network widths (e.g., 2, 4, 8, and 16 neurons) concurrently[13].

The v2 controller employs an oracle-routed regime and introduces a highly rigorous "probation" protocol[13]. When a candidate specialist passes initial validation, it is not immediately merged into the permanent model. Instead, it is published as a reversible probation graph[13]. The framework then waits to collect two further, completely independent decision datasets from the ongoing clinical stream[13]. The exact same gate and weights must independently pass validation on both of these subsequent datasets; if it fails, the controller executes a PROBATION_ROLLBACK, restoring the incumbent graph[13]. This transaction-level probation mirrors the rigorous post-market surveillance and staged rollout required by medical regulators for adaptive software[13].

Empirical results from MORPH v2 validated the structural search approach. Under a continual-oracle regime across 10 paired seeds, the V2_DYNAMIC model (which dynamically searched for optimal widths) yielded a mean balanced MSE of 0.076599 utilizing an observation-averaged 37.47 parameters[13]. In benchmarking against fixed-capacity controls, the

dynamic model required substantially fewer resident parameters during the early phases of learning compared to a static three-expert network (STATIC_3_EXPERT), which consumed 68.00 parameters regardless of task complexity[13].

Continual-Oracle Model Variant	Mean Balanced MSE	Mean Product MSE	Mean Resident Params	Mean Search MS
V2_DYNAMIC	0.076599	0.11474	37.47	69.596
MLP_CONTINUAL	0.091692	0.066409	48.00	10.873
STATIC_3_EXPERT	0.040804	0.030369	68.00	6.024
NO_GROW	1.3825	1.551	4.00	0.412

Table 2: Aggregated empirical outcomes for MORPH v2 continual-oracle benchmarks. Note the efficiency of V2_DYNAMIC in balancing error metrics against parameter footprint[13].

MORPH v3: Polymorphic Numeric Exploration and Edge Constraints

The most advanced iteration, MORPH v3, simulates the ultimate deployment reality of heterogeneous hardware at the clinical edge. Medical devices ranging from powerful radiology workstations to low-power wearable biosensors require different mathematical precision and optimization routines[14]. Rather than restricting the model to a single 32-bit floating-point data type (Float32) and standard backpropagation, v3 conducts a massive polymorphic numeric census across 34 native data types (including Float16, BFloat16, Int8, and Int4) and 31 named training modes[14].

The v3 architecture evaluates 2,108 distinct combinations of data types, network topologies, and learning rates per census[14]. It tests advanced execution modes such as MeshBP (which utilizes an active single-site grid with step Forward/Backward passes) and StepBP (which uses a real FIFO tick line for gradient updates), explicitly rejecting unsupported modes like Adversarial or Memory networks that lack the necessary structural contracts[14].

Crucially, MORPH v3 introduces stringent tracking of "native weight bytes" and "time-weighted resident parameters"[14]. In clinical edge deployments (e.g., Point-of-Care Ultrasound or PoCUS), memory footprint and inference latency are primary deployment blockers[37]. By tracking the exact native payload bytes (excluding the overhead of the Go runtime heap) and quantifying how long candidate shadows reside in memory during the search phase, v3 provides a realistic assessment of the true hardware cost of continual learning[14]. The empirical results

demonstrate that while dynamic growth saves active parameters, the extensive search process imposes a high transient computational cost, necessitating careful configuration of max-shadow-update budgets to prevent battery drain on mobile medical devices[14].

Managing Multimodal Clinical Data via Mixture of Experts (MoE)

While dynamic architectures like MORPH solve the temporal challenge of sequentially arriving tasks, the spatial challenge of clinical data involves processing highly heterogeneous, multimodal inputs simultaneously. Modern healthcare generates fragmented data streams for a single patient, routinely combining structured Electronic Health Records (EHR), unstructured clinical notes, genomic sequences, and high-dimensional medical imaging[15].

Traditional deep learning models attempt to concatenate these diverse modalities into a single, dense representation space, frequently suffering from negative knowledge transfer—where the optimization of a dominant modality (e.g., large high-resolution MRI arrays) drowns out the subtle but critical signals in a minority modality (e.g., a specific genetic mutation code)[34].

To resolve this, advanced clinical AI is transitioning to Mixture of Experts (MoE) architectures[15]. In a modality-driven MoE framework, such as the proposed MoE-Health architecture, the neural network is explicitly modularized[42]. Distinct sub-networks (the "experts") are independently optimized for specific data types[42]. For instance, a vision-transformer expert is allocated to process radiological scans, while a Large Language Model (LLM) expert, perhaps initialized on biomedical corpora like ClinicalBERT or BioBERT, is allocated to extract semantic meaning from physician notes[30]. A learnable gating network then acts as a clinical triage system, calculating sparse activation weights to route the patient's data only to the most relevant experts, dynamically fusing their outputs into a cohesive clinical prediction[43].

This MoE paradigm is equally critical in specialized fields such as spatial transcriptomics. The GC-MoE (Genomics-Guided Cell-Type-Specific Mixture of Experts) architecture applies MoE routing to high-resolution histological data[44]. By allocating specific experts to distinct histological cell types, the network can model localized, cell-specific gene expression gradients[44]. This prevents the network from being overwhelmed by the macroscopic morphological noise of the surrounding tissue, isolating the physiological dynamics of the microenvironment to uncover subtle disease signatures in solid tumors or neurodegenerative tissues[44].

Federated Continual Learning (FCL) and Edge AI Deployment

The transition of clinical AI from centralized academic research clusters to decentralized, real-world deployment necessitates solving the challenges of data sovereignty and hardware limitations.

Overcoming Dual Catastrophic Forgetting in FCL

Federated Learning (FL) enables multiple institutions to train global machine learning models collaboratively by exchanging encrypted parameter gradients with a central server, ensuring that raw, protected health information (PHI) never leaves the local hospital's firewall[9]. However, standard FL algorithms (like FedAvg) assume that the data residing at each hospital is static and identically distributed[34].

Federated Continual Learning (FCL) sits at the highly complex intersection of these domains, addressing environments where multi-center data streams continuously and heterogeneously[34]. FCL faces a compound challenge: models must resist *local* catastrophic forgetting (learning a new local diagnostic task without overwriting prior local knowledge) while simultaneously avoiding *global* catastrophic forgetting (aggregating updates from multiple hospitals without destroying the global model's generalized representations)[34]. Furthermore, non-IID data across different hospitals induces "client drift," where local models over-optimize for their specific demographic biases, skewing the global aggregation[34].

Advanced FCL frameworks, such as DP-FedEPC, integrate Elastic Weight Consolidation (EWC) parameter penalties with prototype rehearsal[9]. Recognizing that storing a replay buffer of raw patient X-rays violates privacy protocols, DP-FedEPC mandates that each client hospital maintains only a highly compressed memory of per-class latent prototypes and their corresponding logits[9]. These prototypes act as a compressed replay buffer that preserves vital decision boundaries[9]. To aggregate these updates securely, the system utilizes Client-Side Differential Privacy (DP-SGD) to inject mathematically calibrated noise into the gradients, providing formal guarantees that the inclusion of any single patient's radiograph cannot be reverse-engineered or memorized by the global model[9].

Edge AI and Point-of-Care Testing (POCT)

The deployment of CL models faces severe physical constraints when moving from centralized hospital servers to edge devices in Point-of-Care Testing (POCT). POCT devices, such as portable ultrasound probes, smart wearable biosensors, and smartphone-integrated diagnostic tools, enable rapid, on-site diagnostics for heart failure, continuous glucose monitoring, and remote infectious disease screening[37].

At the clinical edge, inference latency, power consumption, and memory overhead dominate deployment feasibility[37]. Backpropagation—the core mechanism of continuous learning—is highly memory-intensive and computationally expensive for the standard microcontrollers (MCUs) embedded in these devices[16]. To address these hardware bottlenecks, researchers are shifting away from traditional Convolutional Neural Networks (CNNs) toward novel architectures like Kolmogorov-Arnold Networks (KANs) for edge-based Human Activity Recognition (HAR) and bio-signal processing[16]. KANs feature localized plasticity and global stability, allowing specific mathematical splines within the network to be updated in response to concept drift (e.g., a shift in wearable sensor placement or a change in user biometrics) without triggering the massive, network-wide gradient computations that drain mobile battery

life[16].

Regulatory Paradigms for Adaptive Clinical AI

The fundamental legal and ethical incompatibility between continuous learning systems and traditional medical device regulation is that historical frameworks were engineered exclusively for static hardware. Traditionally, any significant modification to a medical device's design, algorithm, or clinical performance required a completely new premarket submission—a regulatory bottleneck that can take months of review and cost hundreds of thousands of dollars per iteration[49]. To resolve this, international regulatory bodies are actively redesigning their legal frameworks to govern the *lifecycle* of the software rather than the static product itself[51].

FDA 21 CFR 892.2055 and the Predetermined Change Control Plan (PCCP)

The United States Food and Drug Administration (FDA) has pioneered the regulation of adaptive AI through the codification of the Predetermined Change Control Plan (PCCP) pathway. Empowered by Section 515C of the FD&C Act (added via the Food and Drug Omnibus Reform Act of 2022), the FDA established a new regulatory device category on June 17, 2026, under 21 CFR 892.2055: "Radiological machine learning-based quantitative imaging software with predetermined change control plan"[17].

Crucially, this rule places these continuously learning devices into Class II (subject to 510(k) clearance and special controls) rather than the highly burdensome Class III Premarket Approval (PMA) pathway that historically captured unknown AI algorithms[17]. The PCCP is a prospective regulatory agreement filed during the initial 510(k) or De Novo submission. It defines the exact boundaries within which an AI algorithm is permitted to autonomously or manually update in the post-market environment without requiring a new FDA submission[49].

The finalized FDA guidance, issued in December 2024, dictates three mandatory components for a valid PCCP:

1. **Description of Modifications:** A highly specific, bounded list of intended future changes (e.g., retraining on new demographic data, expanding compatibility to new CT scanner models, or updating quantitative output thresholds). Changes must strictly maintain the device's original intended use; sweeping expansions of clinical indications are prohibited under a PCCP[10].
2. **Modification Protocol:** The rigorous methodological framework detailing how modifications will be executed. This encompasses four critical sub-components: data management practices (annotation methods, data segregation, avoidance of data leakage), retraining practices, pre-defined performance evaluation (acceptance criteria and specific statistical metrics like Dice scores evaluated on independent, sequestered test sets), and the update procedures (including fail-safes and criteria to roll back a degraded model)[10].
3. **Impact Assessment:** A comprehensive risk-benefit analysis detailing how the

cumulative implementation of all proposed modifications might introduce new hazards or unmitigated risks, cross-referenced with the mitigation strategies in the Modification Protocol via a traceability matrix[10].

The establishment of 21 CFR 892.2055 mandates "special controls," which legally require developers to maintain transparent version histories, execute robust cybersecurity measures, and continuously monitor Real-World Performance (RWP) for signs of data drift or algorithmic bias[17]. A recent cross-sectional scoping review of FDA-cleared radiology AI devices revealed that following the issuance of the final guidance, PCCP adoption surged, capturing 91.9% of AI/ML-specific PCCPs in radiology[57]. However, the study noted that public regulatory summaries frequently lacked transparent documentation of continuous monitoring metrics and predefined drift triggers, indicating that industry adherence to the transparency tenets of Good Machine Learning Practice (GMLP) is still maturing[57].

The Australian TGA Regulatory Ecosystem

Concurrently, the Australian Therapeutic Goods Administration (TGA) has modernized its approach to Software as a Medical Device (SaMD) through its Essential Principles framework. The TGA evaluates AI software objectively based on its intended purpose—derived not just from the developer's intent, but from technical documentation, marketing claims, and instructions for use[11]. Software intended for the diagnosis, prevention, monitoring, prediction, prognosis, or treatment of disease is strictly classified and regulated as a medical device[58].

The TGA's guidance heavily emphasizes alignment with internationally recognized standards to satisfy the Essential Principles. These include IEC 62304 for rigorous software lifecycle processes, ISO 14971 for comprehensive risk management, IEC 62366-1 for usability engineering, and ISO 13485 for quality management systems[18]. Furthermore, the TGA requires adherence to the IEC 80001 series and ISO/IEC 29147/30111 for cybersecurity and vulnerability disclosure, recognizing that adaptive AI models represent unique cybersecurity attack vectors[18].

A critical distinction in the TGA's enforcement posture relates to "scope creep" and "off-label" use, particularly concerning generative AI and Large Language Models (LLMs)[11]. If a manufacturer becomes aware that a general-purpose LLM or administrative AI tool is being utilized off-label by clinicians for specific diagnostic triage or treatment planning, the TGA expects active and immediate intervention[11]. The manufacturer must either implement hard technical controls to prevent the off-label diagnostic use or halt supply and seek formal ARTG (Australian Register of Therapeutic Goods) inclusion for the expanded medical purpose[11].

Furthermore, the TGA maintains a firm stance on data provenance, stating that while synthetic data can augment training pipelines, it cannot solely replace real-world clinical validation data for the purposes of demonstrating safety and performance[11].

Clinical Translation and the Melbourne Biomedical AI Ecosystem

The theoretical advancements in algorithmic continual learning and regulatory reform are ultimately validated through clinical translation. This translation is heavily dependent on highly integrated research precincts that co-locate computational scientists, clinical oncologists, and regulatory experts. The Melbourne Biomedical AI ecosystem serves as a premier global example of this integration, coordinating the activities of leading institutions including the University of Melbourne, the Peter MacCallum Cancer Centre, St Vincent's Hospital, the Walter and Eliza Hall Institute of Medical Research (WEHI), and Monash University[19].

Multi-omics and Spatial Biology

A major focus of the Melbourne precinct is the synthesis of spatial biology and AI-driven multi-omics. Multi-omics combines data from different biological layers—genomics, proteomics, and transcriptomics—to build a holistic, systems-biology picture of cellular microenvironments[64]. Because acquiring multiple high-resolution biological datasets from a single tissue sample is prohibitively expensive and technically challenging, researchers at Monash University's Biomedicine Discovery Institute developed NicheTrans[64].

NicheTrans is a spatially-aware, transformer-based multi-modal deep learning framework designed to learn the complex, non-linear relationships between different molecular data types[45]. Utilizing spatial context, it can accurately impute and generate missing multi-omics information from much cheaper single-omics measurements[45]. During extensive validation on neurological and oncological datasets, NicheTrans successfully translated protein markers into spatial landmarks, uncovering amyloid-beta associated cell states in Alzheimer's disease and mapping complex glial cell organizations that could not be detected through single-omics analysis alone[45]. By lowering the financial and technical barriers to high-end spatial biology, tools like NicheTrans democratize advanced precision medicine research, enabling laboratories globally to generate high-quality multi-omic insights from affordable diagnostic hardware[64].

Precision Oncology and Genomic Profiling

The integration of AI and comprehensive genomic profiling is actively reshaping clinical oncology pathways across Australia. The Omico consortium, partnering heavily with Melbourne institutions, drives the Molecular Screening and Therapeutics (MoST) program and its critical subprogram, ASPiRATION[62].

The ASPiRATION trial, operating extensively through the Peter MacCallum Cancer Centre, conducts comprehensive genomic profiling (CGP) on patients with newly diagnosed metastatic non-squamous non-small cell lung cancer (mNSCLC)[62]. Traditional pathology sequences genes sequentially; the AI-assisted CGP pipeline analyzes hundreds of genetic biomarkers simultaneously[66]. After a patient's tumor is profiled, the system generates a comprehensive report matching the patient's specific genetic mutations to targeted therapeutic clinical trials[66]. By tracking these interventions, ASPiRATION is capturing vital health-economic real-world evidence to determine whether molecular screening can be feasibly integrated into standard Australian clinical practice, significantly altering the standard of care for lung cancer patients[66].

Advanced Diagnostic Imaging and Adaptive Trials

The clinical translation of AI extends deeply into radiological imaging and trial design. The PRIMARY2 phase III clinical trial, led by the Peter MacCallum Cancer Centre in Melbourne and St Vincent's Hospital in Sydney, leverages advanced diagnostic imaging to optimize prostate cancer biopsy decisions[20]. The trial utilizes Prostate-Specific Membrane Antigen (PSMA) PET/CT scans for patients exhibiting high clinical risk (e.g., strong family history) but who have returned normal MRI results[20]. The radioactive tracer makes aggressive prostate cancer cells highly visible on the scan. This allows clinicians to definitively rule out biopsies for patients with slow-growing or absent disease, saving them from invasive surgical procedures while accurately triaging high-risk individuals into immediate care[20].

Furthermore, institutions like Deakin University are pioneering "adaptive experimental design" utilizing machine learning to drastically accelerate clinical trial timelines[67]. In extensive studies addressing psychological distress, Deakin's Applied Artificial Intelligence Initiative utilizes supercomputers, probability theory, and "lean data" to analyze real-time patient feedback across multiple intervention variables, such as mindfulness, sleep, and diet programs[67]. Instead of waiting years to conclude a static, double-blind clinical trial, this adaptive algorithm continuously refines the treatment protocol to identify the most effective interventions for specific demographics in a matter of months[67]. This adaptive design represents a clinical mirror to the computational continuous learning paradigm, demonstrating how AI can push the entire healthcare apparatus toward true, real-time personalized medicine[67].

Conclusion

The evolution of clinical artificial intelligence has moved definitively past the era of locked, monolithic models. The future of medical diagnostics relies on algorithmic architectures capable of continuous adaptation. By employing dynamic neural network expansion, prototype rehearsal, and Mixture of Experts routing, computer scientists are successfully navigating the complexities of multimodal clinical data while mathematically mitigating the risks of catastrophic forgetting. As demonstrated by frameworks ranging from edge-deployed Kolmogorov-Arnold Networks for point-of-care ultrasound to federated, privacy-preserving networks aggregating data across global hospitals, the technological solutions for spatial and temporal data drift are rapidly maturing.

Crucially, the regulatory landscape has adapted in tandem to support these innovations safely. The codification of 21 CFR 892.2055 and the widespread adoption of the Predetermined Change Control Plan by the FDA, alongside stringent lifecycle governance and standards enforcement by the Australian TGA, establish a robust legal mechanism for adaptive AI. By shifting the regulatory focus from static product clearance to proactive data governance, rigorous modification protocols, and continuous post-market real-world performance monitoring, regulators ensure that algorithms can safely learn from their environments. Translated through highly collaborative research ecosystems like Melbourne Biomedical AI, these adaptive paradigms are lowering the barriers to spatial multi-omics, accelerating precision oncology trials, and ultimately promising a healthcare system that is as dynamic and

resilient as the human biology it seeks to treat.

Works cited

1. Continual Learning in Medical Imaging from Theory to Practice - arXiv, https://arxiv.org/html/2405.13482v1
2. Detecting and Remediating Harmful Data Shifts for the Responsible, https://pmc.ncbi.nlm.nih.gov/articles/PMC12138723/
3. Continual learning in medical devices: FDA's action plan and beyond, https://www.researchgate.net/publication/351163417_Continual_learning_in_medical_devices_FDA's_action_plan_and_beyond
4. Continual Learning for medical image classification. - ResearchGate, https://www.researchgate.net/publication/386448862_Continual_Learning_for_medical_image_classification
5. Continual Learning in Medical Imaging Analysis - arXiv, https://arxiv.org/html/2312.17004v1
6. Should the Use of Adaptive Machine Learning Systems in Medicine, https://www.tandfonline.com/doi/full/10.1080/15265161.2024.2337429
7. Towards continual learning in medical imaging - arXiv, https://arxiv.org/html/1811.02496v1
8. Artificial Intelligence-Enabled Medical Devices - FDA, https://www.fda.gov/medical-devices/digital-health-center-excellence/artificial-intelligence-enabled-medical-devices
9. Federated Continual Learning for Privacy-Preserving Hospital, https://arxiv.org/pdf/2601.06742
10. FDA Predetermined Change Control Plan (PCCP) - CASRAI, https://casrai.org/guides/fda-predetermined-change-control-plan
11. TGA's AI guidance means for AI and medical devices - Allens, https://www.allens.com.au/insights-news/insights/2026/04/same-rules-smarter-tools-what-the-tgas-ai-guidance-means-for-ai-and-medical-devices/
12. v1.pdf
13. v2.pdf
14. v3.pdf
15. MoE-Health: Expert Framework for Healthcare | PDF - Scribd, https://www.scribd.com/document/910300250/MoE-Health-A-Mixture-of-Experts-Framework-for-Robust-Multimodal-Healthcare-Prediction
16. (PDF) COOL: continual online on-device learning for human activity, https://www.researchgate.net/publication/403921923_COOL_continual_online_on-device_learning_for_human_activity_recognition_enhanced_by_KANs
17. FDA's New PCCP Rule for AI Imaging Software: What Changed, https://encord.com/blog/fda-pccp-rule/
18. Australia TGA 2026 Guidance on AI Medical Software Regulation, https://www.pureglobal.com/news/australia-tga-2026-guidance-on-ai-medical-software-regulation
19. Melbourne Biomedical AI, https://machaustralia.org/mbai/

20. Scan that makes prostate cancer cells glow could cut need for biopsies, https://oncologynews.com.au/latest-news/eau-2026-scan-that-makes-prostate-cancer-cells-glow-could-cut-need-for-biopsies/
21. Find Cancer Clinical Trials - Victorian Cancer Trials Link, https://trials.cancervic.org.au/search/
22. Modern Continual Learning with Foundation Models, Evaluation, https://www.preprints.org/manuscript/202605.1234
23. Regulating radiology AI medical devices that evolve in their lifecycle, https://arxiv.org/html/2412.20498v3
24. Dynamic memory to alleviate catastrophic forgetting in continual, https://pmc.ncbi.nlm.nih.gov/articles/PMC8479083/
25. A Novel Continual Learning-based Mobile Application for Screening, https://arxiv.org/html/2410.12589v1
26. Freezing the Physiological Encoder for Auditable Model Updating, https://arxiv.org/html/2607.19020v2
27. Benchmarking MedMNIST dataset on real quantum hardware - PMC, https://pmc.ncbi.nlm.nih.gov/articles/PMC12992601/
28. MedMNIST/README.md at main - GitHub, https://github.com/MedMNIST/MedMNIST/blob/main/README.md
29. MedMNIST v2: A Large-Scale Lightweight Benchmark for 2D and 3D, https://www.researchgate.net/publication/355731017_MedMNIST_v2_A_Large-Scale_Lightweight_Benchmark_for_2D_and_3D_Biomedical_Image_Classification
30. Continually-Adaptive Representation Learning Framework for Time, https://pmc.ncbi.nlm.nih.gov/articles/PMC11920885/
31. Continual Learning in Medical Imaging: A Survey and Practical, https://dl.acm.org/doi/full/10.1145/3785663
32. Continual Learning for Histopathology Image Classification in Class, https://www.mdpi.com/2075-4418/16/11/1711
33. Continual Learning for Histopathology Image Classification in Class, https://pmc.ncbi.nlm.nih.gov/articles/PMC13256828/
34. Federated Continual Learning: Concepts, Challenges, and Solutions, https://arxiv.org/html/2502.07059v2
35. Federated Continual Learning for Privacy-Preserving, Reliable and, https://www.sciltp.com/journals/tai/articles/2603003269
36. Guiding Principles for Predetermined Change Control Plans ... - FDA, https://www.fda.gov/medical-devices/artificial-intelligence-enabled-medical-devices/predetermined-change-control-plans-machine-learning-enabled-medical-devices-guiding-principles
37. Attention-Based Medical Image Analysis: Architectures, Applications, https://www.sciopen.com/article/10.32604/cmc.2026.075316
38. Edge-Optimized Embedded System for Low-Latency Fault Detection, https://pdfs.semanticscholar.org/a633/8d1586bc24a632fd3436fab1c385a37eba44.pdf
39. Intelligent Medical Imaging Systems for Rapid ... - NEXUS PRESS, https://www.nexus-press.com/journal/AMSR/article/2471/view/1

40. MoE-Health: A Mixture of Experts Framework for Robust Multimodal, https://www.researchgate.net/publication/395125509_MoE-Health_A_Mixture_of_Experts_Framework_for_Robust_Multimodal_Healthcare_Prediction
41. Clinical-MoE: Mixture-of-Experts Foundation Models for Electronic, https://www.researchgate.net/publication/406243603_Clinical-MoE_Mixture-of-Experts_Foundation_Models_for_Electronic_Health_Records
42. Modality-driven Mixture-of-Experts - Emergent Mind, https://www.emergentmind.com/topics/modality-driven-mixture-of-experts-moe
43. QA-MoE: Quality-Aware and Stable Multimodal Mixture-of-Experts, https://ijcai.org/proceedings/2026/766
44. GC-MoE: Genomics-Guided Cell-Type-Specific Mixture of Experts, https://arxiv.org/html/2606.02424v1
45. (PDF) NicheTrans: Spatial-aware Cross-omics Translation, https://www.researchgate.net/publication/386600645_NicheTrans_Spatial-aware_Cross-omics_Translation
46. One-Shot Federated Class-Incremental Learning for Medical ... - IJCAI, https://ijcai.org/proceedings/2026/747
47. Jointly exploring client drift and catastrophic forgetting in dynamic, https://pmc.ncbi.nlm.nih.gov/articles/PMC11836390/
48. AI-Reinforced Wearable Sensors and Intelligent Point-of-Care Tests, https://pmc.ncbi.nlm.nih.gov/articles/PMC11595538/
49. Predetermined Change Control Plan (PCCP) for SaMD: FDA Guide, https://operonstrategist.com/predetermined-change-control-plan-pccp-samd/
50. FDA PCCP for AI/ML Medical Devices — 510(k) / De Novo Submission, https://aws.amazon.com/marketplace/pp/prodview-pp6mcaxo4h2og
51. An Overview of Continuous Learning Artificial Intelligence-Enabled, https://canjhealthtechnol.ca/index.php/cjht/article/download/eh0102/704?inline=1
52. FDA PCCP Guidance: Predetermined Change Control Plans, https://superintelligencecompliance.ai/guides/fda-pccp-ai-medical-devices
53. Classification of the Radiological Machine Learning-Based, https://www.federalregister.gov/documents/2026/06/17/2026-12166/medical-devices-radiology-devices-classification-of-the-radiological-machine-learning-based
54. FDA PCCP for AI Medical Devices | Change Control Plan Guide, https://www.tactionsoft.com/blog/fda-pccp-ai-medical-devices/
55. Understanding Predetermined Change Control Plans (PCCP) for, https://www.sushvin.com/understanding-predetermined-change-control-plans-for-medical-devices.html
56. Navigating FDA AI Guidance for Medical Devices: A Practical Guide, https://www.jamasoftware.com/blog/navigating-fda-ai-guidance-for-medical-devices-a-practical-guide/
57. Predetermined Change Control Plan Adoption and Documentation, https://pubs.rsna.org/doi/10.1148/ryai.260385
58. Healthcare software: TGA clarifies when AI will be... | Clayton Utz, https://www.claytonutz.com/insights/2026/july/healthcare-software-tga-clarifies-

when-ai-will-be-regulated-as-a-medical-device
59. Understanding how we regulate software-based medical devices, https://www.tga.gov.au/resources/guidance/understanding-how-we-regulate-software-based-medical-devices
60. Standards for software-based medical devices, https://www.tga.gov.au/products/medical-devices/software-and-artificial-intelligence-ai/overview/standards-software-based-medical-devices
61. Software and AI medical device compliance, https://www.tga.gov.au/products/medical-devices/software-and-artificial-intelligence-ai/manufacturing/artificial-intelligence-ai-and-medical-device-software-regulation/software-and-ai-medical-device-compliance
62. Outsmarting cancer, together - Omico, http://omico.org.au/wp-content/uploads/2025/10/Omico_Annual-Report_2025_V2.2.pdf
63. MACH-Track - Melbourne Academic Centre for Health, https://machaustralia.org/future-leaders/mach-track/
64. Monash AI-driven biomedical model could make advanced disease, https://www.monash.edu/discovery-institute/news-and-events/news/2026-articles/monash-ai-driven-biomedical-model-could-make-advanced-disease-research-faster,-cheaper-and-more-accessible
65. NicheTrans: Spatial-aware Cross-omics Translation - bioRxiv, https://www.biorxiv.org/content/10.1101/2024.12.05.626986v1.full.pdf
66. ASPiRATION - ANZCTR - Registration, https://anzctr.org.au/Trial/Registration/TrialReview.aspx?id=381073
67. Artificial intelligence providing faster and more personalised mental, https://www.deakin.edu.au/study-research-and-career-stories/research/artificial-intelligence-providing-faster-and-more-personalised-mental-health-treatment