

Memory-Augmented Multi-Omics Learning for Predicting Exercise-Induced Transcriptomic and Splicing Responses in Human Skeletal Muscle

Guangpeng Dai

Department of Electrical Engineering and Computer Science, University of Missouri,
Columbia, MO, USA.

guangpeng.dai@missouri.edu

Brooks L. Young

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.

young1984@buffalo.edu

Abstract

Predicting individualized molecular responses to acute and chronic exercise is a grand challenge that sits at the confluence of systems genomics, artificial intelligence, and precision health. Human skeletal muscle remodeling involves coordinated transcriptomic and alternative splicing programs that unfold over multiple temporal scales, and these programs are modulated by genetic variation, epigenetic state, and environmental context. Traditional machine learning models, constrained by fixed-size context windows and static input representations, are ill-suited for capturing the nonstationary and interdependent dynamics of multi-omics time courses. This paper presents a system-level analysis of memory-augmented neural architectures for learning both transcript abundance and splicing isoform trajectories following exercise stimuli. We examine the structural trade-offs inherent in coupling external memory banks with cross-omics attention mechanisms, discuss the infrastructural requirements for training such models on large-scale human exercise cohorts, and evaluate robustness under data heterogeneity, missingness, and population stratification. Governance challenges including privacy-preserving data sharing, fairness across ancestry groups, algorithmic interpretability, and the sustainability of computationally intensive training pipelines are addressed. Integrating the emerging generation of wearable biosensors and genetically encoded ionic-stress reporters further expands the data landscape and demands anticipatory policy frameworks. By situating memory-augmented multi-omics learning within a socio-technical infrastructure perspective, this work delineates a roadmap for deploying predictive models that are not only accurate but also equitable, interpretable, and operationally viable at scale.

Keywords

memory-augmented neural networks; multi-omics integration; transcriptomics; alternative splicing; skeletal muscle; exercise genomics; fairness; system architecture.

1. Introduction

Human skeletal muscle exhibits remarkable plasticity in its transcriptome and splice isoform landscape in response to endurance, resistance, and high-intensity interval exercise. High-throughput RNA sequencing has revealed thousands of genes whose expression changes within hours of an acute bout, along with widespread shifts in alternative splicing that

generate functionally distinct protein variants [1]. Recent investigations further demonstrate that single-nucleotide polymorphisms modulate both baseline expression levels and the magnitude of exercise-induced transcriptomic and splicing responses when coupled with dietary interventions [3]. Despite this wealth of data, predicting an individual's molecular trajectory following a specific exercise protocol remains an open problem, largely because the regulatory interplay among transcriptional, post-transcriptional, and epigenomic layers operates over long temporal dependencies that extend well beyond a single time point. Exercise perturbs a highly connected system in which the effect of a transcription factor binding event in the early recovery phase may manifest as a splicing change hours later, mediated by chromatin remodeling and metabolic signaling cascades.

Conventional predictive models, ranging from linear mixed-effects frameworks to deep feedforward or recurrent networks, typically process a fixed window of omics features. These models suffer from a fundamental mismatch with the nature of exercise biology: the relevant conditioning context required to predict a future state often resides far back in the time series or in a different molecular layer altogether. Moreover, alternative splicing predictions require an understanding of sequence grammar and regulatory element combinations that span intronic and exonic regions, a problem that has seen major advances through deep learning models such as SpliceAI [10]. However, integrating transcript- and splicing-level predictions into a single coherent architecture that also accommodates genetic variation and preceding exercise history demands a new class of models with explicit external memory. Memory-augmented neural networks, originally developed for question answering and one-shot learning tasks, are uniquely positioned to fill this gap by decoupling the model's parametric knowledge from the dynamic storage of episode-specific information [5]. This paper adopts an interdisciplinary systems perspective to articulate the design, deployment, and governance of such memory-augmented multi-omics learners. We focus on structural trade-offs, infrastructural implications, robustness, fairness, sustainability, and policy rather than on a narrow algorithmic contribution.

2. The Multi-Omics Landscape of Exercise-Induced Skeletal Muscle Remodeling

The transcriptomic response to a single bout of exercise is remarkably transient, with peak changes occurring within a few hours and returning toward baseline within 24 hours, while training-induced adaptations accumulate over weeks through repeated stimuli. Therefore, the predictive challenge inherently spans multiple timescales, requiring a system that can maintain what is effectively a compressed memory of past perturbations and the molecular state of the muscle at each intervention point. Alternative splicing adds another layer of complexity: a single gene may produce several isoforms with antagonistic functions, and exercise can shift the ratio between them in a manner that is not simply proportional to the change in overall gene expression [2]. Splicing decisions are influenced by RNA-binding protein concentrations, splice site strengths, and epigenetic marks, all of which evolve during recovery and can be modulated by genetic variants that lie outside canonical splice sites [3].

Multi-omics integration frameworks that jointly analyze transcriptomic, proteomic, and epigenomic data have been developed primarily in the context of cancer and developmental biology [4]. Extending these frameworks to exercise biology requires explicit modeling of the temporal axis, because cross-omics relationships are lagged and context-dependent. For example, an epigenetic modification deposited during an acute bout may prime a gene for faster transcriptional re-activation in a subsequent session, a phenomenon that cannot be captured without a memory mechanism that links episodes. The inclusion of genotype data, as

demonstrated in exercise-diet interaction studies, further increases the state space: a missense variant might alter the half-life of a signaling protein, thereby changing the relaxation timescale of an entire regulatory module [3]. Thus, the system must learn what to remember and for how long, which is precisely the capacity that memory-augmented architectures provide.

3. Memory-Augmented Architectures for Temporal Multi-Omics Learning

Memory-augmented neural networks, including end-to-end memory networks and the differentiable neural computer, equip a controller network with an external memory matrix that can be written to and read from through differentiable attention operations [5]. In the exercise multi-omics context, the controller could be a transformer-based encoder that receives concatenated transcriptomic and splicing features at the current time point, while the external memory stores embeddings of earlier molecular states, exercise type tags, and static genetic features. Each training episode would correspond to a multi-time-point trajectory for a given individual, enabling the model to learn retrieval strategies that condition the current prediction on memories of similar past states. The memory module thus serves as a learned representation of an individual's internal regulatory dynamics, much as memory networks store factual knowledge for question answering.

A critical structural trade-off concerns the sizing of the memory bank. A memory that is too small will collapse diverse histories into indistinguishable representations, undermining personalization, while an excessively large memory incurs quadratic attention costs in the read and write heads. Recent work on memory-aware fine-tuning for long-sequence video object segmentation has demonstrated that segmented memory management and adaptive compression can maintain high fidelity for extended temporal sequences while keeping computation tractable [6]. These insights are directly transferable to multi-omics trajectories, where a temporal sequence may span dozens of time points and thousands of features. The memory can be partitioned into short-term and long-term banks, with consolidation mechanisms akin to hippocampal replay, ensuring that salient exercise-induced perturbations remain retrievable without continuous storage of redundant resting states.

Another architectural consideration is the cross-omics integration strategy. Transcript-level and splicing-level features may be fused either early, by concatenation before the memory write, or late, through separate memory banks and a gated interaction layer. Late fusion enables the model to learn distinct dynamics for each omics layer and then synthesize them only when a joint prediction is needed, such as forecasting the combined transcriptional and isoform response 24 hours after a bout. Enformer and related models have shown that long-range cis-regulatory interactions, spanning hundreds of kilobases, can be captured by transformers with expansive receptive fields, but they operate on static DNA sequences [7]. Augmenting such sequence-based models with dynamic memory extends their capability to integrate the trans-acting environment of the nucleus, including signaling pathway activity and metabolite concentrations, thereby producing predictions that are conditional on physiological history. The controller may also incorporate a multi-head attention mechanism over the memory contents, aligned with the concept of multi-query read that has proven valuable in meta-learning settings.

4. Infrastructure for Scalable Deployment

Translating a memory-augmented multi-omics model from a proof-of-concept trained on a small cohort to a deployable system that can serve millions of users requires careful attention

to data infrastructure, compute orchestration, and compliance with data protection regulations. Large-scale exercise genomics consortia, such as MoTrPAC and the GTEx project, have begun generating deep multi-omics time courses but remain limited in sample size relative to the dimensionality of genotype-by-time interactions [11]. A viable deployment strategy must therefore federate data across multiple institutions, each housing sensitive human omics data. Federated learning with secure aggregation partially addresses privacy, but the recurrent nature of memory read-write operations introduces additional leakage risks because the memory state can implicitly encode information about individual training samples.

Differential privacy has become a standard safeguard in deep learning, bounding the influence of any single data point on the model parameters through calibrated noise injection [8]. However, applying differential privacy to external memory comes with a unique challenge: noise added to memory writes can accumulate over a long exercise trajectory, degrading the signal for clinically relevant but subtle splicing shifts. A practical compromise is to apply privacy accounting at the level of aggregated trajectory summaries rather than at every time step, striking a balance between utility and formal guarantees. Cloud-native infrastructure, employing containerized microservices and Kubernetes-based autoscaling, can handle the bursty training demands that arise when models are periodically updated with new cohort data. Yet such infrastructure carries a substantial carbon footprint, a concern that intersects with the increasing awareness of computational sustainability [16]. Model compression techniques, including memory sparsification and low-rank factorization of the memory matrix, are not merely optimization afterthoughts but structural design choices that directly affect the environmental cost of model development and refresh cycles.

Data integration also requires robust Extract-Transform-Load pipelines that normalize disparate omics modalities, genotype arrays, and phenotype annotations. The FAIR principles—Findable, Accessible, Interoperable, Reusable—should guide the construction of data lakes that can feed the memory-augmented learner. External data sources such as wearable biosensors for heart rate variability, continuous glucose monitoring, and even flexible bioelectronic patches that measure electrolyte concentrations in interstitial fluid are increasingly capable of providing surrogate physiological context between muscle biopsies [13]. These streaming data streams can be encoded as auxiliary write events into the memory bank, tightening the coupling between observable physiology and predicted molecular state.

5. Robustness, Interpretability, and Fairness in Predictive Genomic Models

A model that accurately predicts exercise responses in one cohort may fail dramatically in another due to distribution shift, batch effects, or underrepresentation of certain ancestry groups. Robustness must be engineered at multiple levels: data augmentation through simulation of diverse exercise protocols and nutritional states, invariant representation learning that removes unwanted sources of variation while preserving individual-specific regulatory signals, and uncertainty calibration that quantifies confidence in splicing variant predictions. Interpretability is equally critical because clinicians and exercise physiologists need to understand why a particular isoform switch is predicted for a given individual. Saliency maps over memory read weights can expose which historical states or genetic loci chiefly drive a current prediction, providing a traceable chain of reasoning.

The interpretability of splicing predictions has been advanced by models that directly score the contribution of sequence motifs and their interactions, making it possible to link a predicted splice site change to a specific variant in a splicing regulatory element [10]. When embedded within a memory-augmented framework, such interpretable motif scoring can be

extended to incorporate the temporal context encoded in memory, for instance showing that a variant exerts its effect only after a certain number of exercise bouts or under a specific nutritional regime. This capability is not only scientifically revealing but also essential for building trust with end-users and for meeting regulatory requirements that increasingly demand explainable artificial intelligence in health applications.

Fairness emerges as a systemic concern because the majority of exercise genomics studies have historically been conducted in populations of European descent [9]. A memory-augmented model trained predominantly on one ancestry group will overfit to allele frequency spectra and linkage disequilibrium patterns that are not representative globally, potentially issuing inaccurate or even harmful recommendations for exercise prescription in underrepresented groups. Algorithmic fairness techniques, including reweighting, adversarial debiasing, and multi-task learning with ancestry as an auxiliary output, must be incorporated into the training objective. Fairness audits should be conducted not simply on aggregate prediction error but on performance across strata defined by self-reported ancestry and admixed genomes, paying particular attention to false negative rates for beneficial isoform adaptations that may be differentially distributed. Infrastructure for such audits requires diverse validation cohorts and transparent reporting, both of which demand sustained investment and governance structures that go beyond individual research teams.

6. Governance, Sustainability, and Policy Implications

Deploying a memory-augmented multi-omics predictive system at population scale implicates a complex governance landscape spanning data protection, medical device regulation, algorithmic accountability, and global health equity. In many jurisdictions, software that predicts an individual's physiological response to exercise with sufficient specificity to inform clinical decisions may be classified as a medical device, triggering conformity assessment and post-market surveillance obligations. The dynamic nature of memory-augmented models, which continue to update their internal memory representations based on new user data, blurs the line between a locked algorithm and a continuously learning system, a distinction that regulators such as the U.S. Food and Drug Administration and the European Medicines Agency are actively grappling with. Proposed frameworks for adaptive algorithms suggest that change control protocols, pre-specified performance boundaries, and real-world monitoring dashboards must be integral to the deployment architecture from the outset [14].

Sustainability of the technical infrastructure intersects with governance when considering the energy and resource demands of retraining cycles. The carbon footprint of bioinformatics workloads has been scrutinized, showing that frequent model refreshes on large genomic datasets can rival that of other high-performance computing domains [16]. Policy levers such as green procurement mandates for cloud services, mandatory reporting of energy consumption for health AI applications, and shared benchmarking consortiums can collectively steer the community toward memory compression and efficient attention mechanisms that do not compromise predictive performance. Additionally, the production of wearable biosensors and their integration into multi-omics prediction pipelines generates electronic waste and supply chain dependencies that must be managed under extended producer responsibility legislation.

The potential integration of novel biosensors that report on local ionic stress, pH fluctuations, or metabolic intermediate concentrations in real time adds a further dimension to the data ecosystem [12]. Such genetically encoded sensors, demonstrated initially in model organisms, could eventually be adapted for human use in exercise physiology research through minimally

invasive methods. When such high-resolution physiological streams feed into the memory bank of a predictive model, governance questions arise regarding data ownership, consent for continuous monitoring, and the preventability of adverse outcomes when a model flags an abnormal predicted trajectory. Policy anticipatory exercises, involving ethicists, exercise scientists, regulators, and patient advocacy groups, should begin well before such sensors reach commercial maturity, so that the socio-technical infrastructure can be shaped proactively rather than retrofitted under crisis conditions [17].

7. Conclusion

Memory-augmented learning architectures offer a principled route to capturing the long-range, multi-timescale dependencies that characterize exercise-induced transcriptomic and splicing responses in human skeletal muscle. By decoupling parametric knowledge from episodic storage, these systems can personalize predictions without sacrificing generalizability, provided that the infrastructure supporting them is designed with rigorous attention to robustness, fairness, interpretability, and sustainability. The structural choices—early versus late omics fusion, memory bank sizing, differential-privacy noise scheduling, federated aggregation—are not implementation details; they constitute the core of what makes such a system trustworthy and deployable at scale. Governance frameworks must evolve in parallel, addressing adaptive algorithm regulation, equity across populations, environmental cost, and the integration of emerging biosensor data including ionic-stress monitors. A holistic socio-technical perspective, as advanced in this paper, is essential for translating powerful memory-augmented multi-omics models into equitable and enduring tools for precision exercise medicine.

References

1. Lindholm, M. E., Huss, M., Solnestam, B. W., Kjellqvist, S., Lundeberg, J., & Sundberg, C. J. (2014). The human skeletal muscle transcriptome: sex differences, alternative splicing, and tissue homogeneity assessed with RNA sequencing. *FASEB Journal*, 28(10), 4571–4581.
2. Ost, M., Igual Gil, C., Coleman, V., Keipert, S., Efstathiou, S., Vidic, V., Weyers, M., & Klaus, S. (2021). Exercise induces isoform-specific changes in mRNA expression and alternative splicing in human skeletal muscle. *Physiological Reports*, 9(10), e14882.
3. Wang, W. et al. Impact of polymorphisms on gene expression and splicing in response to exercise and diet-induced weight loss in human skeletal muscle tissues. *Cell Genom* 5, 100951 (2025).
4. Hasin, Y., Seldin, M., & Lusis, A. (2017). Multi-omics approaches to disease. *Genome Biology*, 18(1), 83.
5. Sukhbaatar, S., Szlam, A., Weston, J., & Fergus, R. (2015). End-to-end memory networks. *Advances in Neural Information Processing Systems*, 28, 2440–2448.
6. Li, G., Yuan, H., Chen, S., Hu, Q., Wang, J., & Jiang, K. (2026). MFT: Memory-Aware Fine-Tuning of SAM2 for Efficient Long-Sequence Video Object Segmentation. *IEEE Signal Processing Letters*.
7. Avsec, Ž., Weilert, M., Shrikumar, A., Krueger, S., Alexandari, A., Dalal, K., Fropp, R., McAnany, C., Gagneur, J., Kundaje, A., & Zeitlinger, J. (2021). Effective gene

expression prediction from sequence by integrating long-range interactions. *Nature Methods*, 18(10), 1196–1203.

8. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318.
9. Popejoy, A. B., & Fullerton, S. M. (2016). Genomics is failing on diversity. *Nature*, 538(7624), 161–164.
10. Jaganathan, K., Kyriazopoulou Panagiotopoulou, S., McRae, J. F., Darbandi, S. F., Knowles, D., Li, Y. I., Kosmicki, J. A., Arbelaez, J., Cui, W., Schwartz, G. B., Chow, E. D., Kanterakis, E., Gao, H., Kia, A., Batzoglou, S., Sanders, S. J., & Farh, K. K.-H. (2019). Predicting splicing from primary sequence with deep learning. *Cell*, 176(3), 535–548.
11. GTEx Consortium. (2020). The GTEx Consortium atlas of genetic regulatory effects across human tissues. *Science*, 369(6509), 1318–1330.
12. Ji, Z., Liu, J., Wang, B., Wei, S., Bian, Y., Zeng, W., Chung, C. I., Ma, Z., Zhang, J., Shu, X., & Ma, D. K. (2026). A genetically encoded ionic-stress sensor reveals protons as a sleep driver. *bioRxiv*. <https://doi.org/10.64898/2026.01.27.702131>
13. Someya, T., Bao, Z., & Malliaras, G. G. (2016). The rise of plastic bioelectronics. *Nature*, 540(7633), 379–385.
14. Floridi, L., Cows, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707.
15. Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G., & King, D. (2019). Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine*, 17(1), 195.
16. Lannelongue, L., Grealey, J., Inouye, M., & Bateman, A. (2021). The carbon footprint of bioinformatics. *Molecular Biology and Evolution*, 38(4), 1791–1798.
17. Vayena, E., Blasimme, A., & Cohen, I. G. (2018). Machine learning in health care: addressing ethical challenges. *PLoS Medicine*, 15(11), e1002689.