

Multimodal Recommender Systems for Predicting Anti-Aging Effects of Animal-Derived Bioactive Peptides

Inkeri Hakkinen

Department of Computer Science, University of New Hampshire, Durham, NH, USA.

inker.work@unh.edu

Hans-Friedrich Solzer

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

hansfriedrichsolzer385@oregonstate.edu

April Bean

Department of Computer Science, University of Central Florida, Orlando, FL, USA.

aprilmail@ucf.edu

Abstract

The prediction of anti-aging effects from animal-derived bioactive peptides is an emerging interdisciplinary problem that combines peptide chemistry, geroscience, nutrition, pharmacology, machine learning, and recommender-system engineering. This paper examines how multimodal recommender systems can support the prioritization and personalization of peptide candidates whose effects may influence aging-related biological processes. Rather than treating recommendation as a narrow ranking task, the paper frames it as a socio-technical infrastructure that integrates heterogeneous evidence from peptide sequences, biochemical assays, omics measurements, imaging, textual literature, clinical records, and consumer or clinical preference signals. The analysis emphasizes system-level design trade-offs involving data fusion, modality imbalance, cold-start conditions, interpretability, robustness, fairness, privacy, deployment, and long-term sustainability. It also considers governance mechanisms such as model documentation, dataset provenance, internal auditing, and policy oversight for health-related recommendations. A central argument is that multimodal recommender systems for anti-aging peptides should not be evaluated only by predictive accuracy, but also by their ability to maintain evidential traceability, avoid overclaiming, protect vulnerable users, and adapt to evolving scientific consensus. The paper concludes that responsible deployment requires modular architecture, transparent uncertainty communication, domain-specific validation, and coordinated oversight across research, clinical, regulatory, and commercial contexts. Such systems may become useful discovery aids, but only if they are designed as accountable scientific instruments rather than as opaque optimization engines.

Keywords

multimodal recommender systems; animal-derived bioactive peptides; anti-aging; system architecture; robustness; fairness; governance; sustainability.

1. Introduction

The search for interventions that modulate aging-related decline has expanded from small-molecule pharmacology and genetic manipulation to nutritionally derived bioactive compounds, including peptides obtained from animal sources. Animal-derived bioactive peptides are fragments of proteins that can exhibit antioxidant, anti-inflammatory, immunomodulatory, metabolic, and tissue-repair activities, making them plausible candidates for interventions aimed at aging-related diseases and functional decline [10]. However, the evidentiary landscape is fragmented. Peptide activity depends on sequence, length, charge, hydrophobicity, bioavailability, digestion, transport, target tissue, dose, duration, and host physiology. These variables are distributed across laboratory assays, animal studies, human trials, omics datasets, patents, and regulatory documents. As a result, identifying promising peptides and matching them to particular aging-related contexts is not simply a matter of searching a database. It requires computational systems that can integrate heterogeneous evidence and produce recommendations that are useful, calibrated, and accountable.

Multimodal recommender systems offer a conceptual and technical framework for this problem. Recommender systems were developed to rank items for users under uncertainty, often by combining collaborative signals, content features, and contextual information [4]. In the anti-aging peptide domain, the “user” may be a researcher, a clinician, a formulation scientist, or a consumer, while the “item” may be a peptide candidate, a peptide combination, a delivery strategy, or a study design. The multimodal dimension arises because peptide evidence is not confined to one data type. It includes sequence representations, spectral data, assay outcomes, gene-expression profiles, proteomic measurements, imaging phenotypes, scientific text, clinical histories, and preference or outcome feedback. Deep learning has made such integration more feasible, particularly through representation learning and attention mechanisms that can align information across modalities [1], [2], [3]. Yet technical feasibility does not guarantee scientific validity or social acceptability.

This paper argues that multimodal recommender systems for predicting anti-aging effects of animal-derived bioactive peptides should be understood as large-scale socio-technical infrastructures. Their value depends not only on model architecture but also on data provenance, validation regimes, deployment environments, incentive structures, and governance. The discussion is organized around system-level concerns: how multimodal evidence should be structured; how recommendation objectives should be defined; how architecture and infrastructure shape scalability and sustainability; how robustness and fairness can be assessed; and how policy and institutional oversight can prevent overclaiming in a domain where health, commerce, and scientific uncertainty intersect.

2. Conceptual Foundations and Problem Framing

A recommender system for anti-aging peptides must first define what is being recommended and for whom. In conventional commerce, the objective may be click-through, purchase, or engagement. In biomedical discovery, the objective is more complex. A recommendation might prioritize a peptide for further laboratory testing, suggest a peptide for a specific aging-related phenotype, propose a combination with synergistic potential, or identify a delivery formulation likely to improve bioavailability. These are different tasks with different error costs. A false positive in consumer-facing anti-aging recommendation may lead to wasted spending or unsafe use, while a false negative in discovery may delay a promising research direction. Therefore, the problem framing must specify the decision context, the evidence threshold, and the acceptable level of uncertainty [5], [6].

Animal-derived bioactive peptides add further complexity because their effects are often pleiotropic and context-dependent. A peptide may reduce oxidative stress in one model system but have limited bioavailability in humans, or it may modulate inflammatory pathways in a way that is beneficial in one tissue and harmful in another. Aging itself is not a single disease but a collection of interrelated processes, including genomic instability, telomere attrition, epigenetic alterations, loss of proteostasis, mitochondrial dysfunction, cellular senescence, stem cell exhaustion, and altered intercellular communication [18]. A recommender system that ignores this multidimensionality may collapse diverse outcomes into a single anti-aging score, thereby losing the nuance needed for scientific and clinical decision-making. A better framing treats anti-aging effects as a profile of outcomes across biological domains, populations, doses, and time horizons.

The recommender paradigm also introduces questions about feedback loops. If researchers repeatedly test peptides recommended by the system, the resulting data may become biased toward the model's prior beliefs. Over time, the system may reinforce dominant hypotheses and underexplore novel peptide classes or underrepresented animal sources. This is a well-known challenge in machine learning systems, where technical debt and feedback dynamics can degrade long-term performance [12]. In biomedical discovery, the consequence is not merely lower accuracy but a narrowing of the scientific search space. A responsible system must therefore incorporate exploration, uncertainty sampling, and diversity-aware ranking, not only exploitation of predicted efficacy.

3. Multimodal Data Architecture for Peptide Effect Prediction

The multimodal data architecture is the foundation of any prediction or recommendation pipeline. Peptide sequence data provide a symbolic representation that can be transformed into numerical features through learned embeddings. Biochemical and biophysical assay data provide direct measurements of activity, stability, toxicity, and binding. Omics data, including transcriptomics, proteomics, and metabolomics, describe system-level responses to peptide exposure. Imaging data can capture morphological or functional changes in cells or tissues. Textual data from scientific literature, patents, and regulatory filings provide evidence about prior findings, mechanisms, and safety signals. Clinical data, when available, link interventions to human outcomes, although privacy constraints and sparse coverage limit their use. Integrating these modalities requires alignment across different scales, vocabularies, and missingness patterns [3], [8].

Graph-based approaches are particularly relevant because biomedical knowledge is relational. Peptides interact with receptors, enzymes, pathways, cell types, and disease phenotypes. Graph neural networks can propagate information across these relationships and learn representations that capture higher-order structure [8], [9]. A multimodal recommender system might combine sequence embeddings, assay features, omics signatures, and graph-based knowledge into a shared latent space. However, the architecture must handle modality imbalance. Sequence data are abundant and standardized, while human clinical evidence for animal-derived peptides is scarce and heterogeneous. If the model is trained primarily on sequence and assay data, it may appear accurate while lacking external validity. A robust architecture should therefore include modality-specific encoders, gated fusion mechanisms, and explicit uncertainty estimates that reflect data density and quality.

Data provenance is as important as data fusion. Datasets differ in species, preparation method, dosage, endpoint definitions, and reporting standards. Without standardized metadata, the model may learn spurious associations. Documentation frameworks such as datasheets for

datasets and model cards for model reporting are useful because they require developers to describe collection processes, intended uses, limitations, and ethical considerations [13], [14]. In the peptide domain, provenance should include animal source, extraction method, purity, sequence verification, assay protocol, model organism, sample size, and statistical treatment. This information should be attached to data points and surfaced in recommendations, allowing users to judge whether a suggested peptide is supported by relevant evidence or by distant analogies.

4. Recommender Systems as Decision Support in Anti-Aging Discovery

Recommender systems can function as decision-support tools that rank peptide candidates for experimental validation. In this role, they resemble drug repurposing and virtual screening systems more than consumer recommendation engines. The system may learn from historical experiments, published associations, and expert judgments. Neural collaborative filtering and graph collaborative filtering can model interactions between peptides, targets, and contexts [6], [9]. Optimization methods such as Bayesian personalized ranking provide ways to learn from implicit feedback, such as which peptides were selected for follow-up studies [7]. However, the objective function must be carefully designed. If the system optimizes only for predicted efficacy, it may recommend compounds with high potency but poor safety, stability, or bioavailability. A multi-objective approach should balance efficacy, uncertainty, novelty, safety, and feasibility.

Personalization introduces additional ethical and scientific tensions. In a clinical or nutritional context, personalization may mean tailoring recommendations to an individual's age, sex, comorbidities, medications, genetic profile, microbiome, and lifestyle. Such tailoring could improve relevance, but it also raises privacy concerns and may create false precision when evidence is weak. In a research context, personalization may mean tailoring recommendations to a laboratory's capabilities, prior data, or disease model. This form of personalization is less ethically fraught but can still create filter bubbles that narrow scientific exploration. A well-designed recommender should distinguish between personalization that is supported by evidence and personalization that is merely an artifact of the model's training distribution [11], [12].

The interpretability of recommendations is essential. Scientists and clinicians need to know why a peptide was recommended. Was it due to sequence similarity, pathway enrichment, prior experimental success, or textual association? A recommendation without explanation cannot be audited or trusted. Post-hoc explanation methods can help, but they are often unstable and may not reflect the model's true reasoning. A more robust approach is to design architectures that expose intermediate evidence, such as attention over modalities, pathway-level scores, and provenance links. In this sense, the recommender system should not only answer which peptide to consider but also provide a traceable rationale that can be challenged, replicated, or refined.

5. System Architecture and Deployment Trade-Offs

Deploying a multimodal recommender for anti-aging peptides involves architectural choices that affect scalability, latency, maintainability, and sustainability. A centralized architecture can simplify training and governance but may concentrate sensitive data and create single points of failure. A federated or distributed architecture can keep clinical and proprietary data local, but it complicates model updates, evaluation, and fairness auditing. In practice, hybrid architectures may be necessary. Public data and literature-derived features can be processed

centrally, while private omics or clinical data remain within institutional boundaries. Federated learning and secure computation can support collaborative model improvement, but they add technical complexity and may reduce transparency unless accompanied by strong documentation and audit mechanisms [15].

Deployment environments also shape design. A research-facing system may prioritize flexibility, reproducibility, and batch processing, while a clinical decision-support system may require real-time inference, audit trails, and integration with electronic health records. A consumer-facing application may prioritize usability, but it also faces greater risk of misinterpretation. Because anti-aging is a domain with strong commercial incentives and weak regulatory boundaries in some jurisdictions, consumer-facing recommendations can easily slide from scientific hypothesis to marketing claim. A responsible deployment strategy should therefore separate exploratory recommendation from medical advice and clearly label evidence levels. Regulatory frameworks for artificial intelligence in health are still evolving, but principles such as transparency, accountability, safety, and human oversight are widely endorsed [16].

Maintenance is another system-level concern. Multimodal models depend on pipelines for data ingestion, preprocessing, feature extraction, model training, evaluation, and monitoring. These pipelines accumulate technical debt when they are tightly coupled, poorly documented, or dependent on fragile data sources [12]. In the peptide domain, new studies appear continuously, assay protocols change, and nomenclature is inconsistent. A sustainable system must support incremental updates without destabilizing prior predictions. Versioning of data, models, and evaluation metrics is essential. Without version control, a recommendation made today may be impossible to reproduce tomorrow, undermining scientific accountability.

6. Robustness, Fairness, and Governance

Robustness refers to a system's ability to maintain reliable performance under distribution shift, noisy inputs, missing modalities, and adversarial manipulation. In anti-aging peptide prediction, distribution shift is pervasive. A model trained on rodent studies may perform poorly on human data. A model trained on one peptide family may fail on another. A model trained before a new assay standard may misinterpret older data. Robustness testing should therefore include external validation across species, preparation methods, and endpoint definitions. It should also test for sensitivity to missing modalities, because real-world queries will rarely have complete multimodal evidence. Uncertainty quantification is not optional; it is a core safety feature. Users should know when the system is extrapolating beyond its training distribution [13], [17].

Fairness is equally complex. In consumer health, unfairness may arise if recommendations systematically favor products from certain regions, sponsors, or demographic groups. In clinical research, unfairness may arise if training data underrepresent older adults, women, or diverse ancestral populations. In discovery, unfairness may mean that peptides from underrepresented animal sources or traditional food systems are ignored. Fairness cannot be reduced to a single metric. It requires participatory design, disaggregated evaluation, and ongoing monitoring. Model cards and datasheets can document known disparities, but they do not by themselves correct them [13], [14]. Governance mechanisms must include channels for affected communities to raise concerns and for independent auditors to inspect system behavior.

Governance also involves accountability for downstream effects. If a recommender system suggests a peptide that is later marketed with exaggerated anti-aging claims, who is responsible? The model developer, the data provider, the platform operator, or the marketer? Internal algorithmic auditing frameworks can help organizations trace decisions, assign responsibility, and remediate harms [15]. Ethical frameworks for artificial intelligence emphasize beneficence, non-maleficence, autonomy, justice, and explicability [16]. In the anti-aging domain, these principles must be interpreted in light of commercial pressure. The promise of longevity can motivate healthy behavior, but it can also exploit fear of aging and sell unproven interventions. A governance regime should therefore require evidence grading, conflict-of-interest disclosure, and restrictions on claims that exceed the system's validated scope.

7. Infrastructure, Sustainability, and Policy Implications

The infrastructure required for multimodal peptide recommendation includes data repositories, compute clusters, model registries, application programming interfaces, identity management, and monitoring systems. These components are not neutral. They determine who can contribute data, who can query models, and who benefits from predictions. Open science infrastructures can democratize access, but they also raise privacy and intellectual property concerns. Proprietary infrastructures may protect commercial interests but reduce reproducibility. A sustainable approach may combine open benchmarks, standardized metadata, and tiered access controls. Public investment in curated peptide and aging datasets could reduce duplication and improve comparability across studies.

Sustainability also has an environmental dimension. Training large multimodal models consumes energy and water, and continuous retraining can increase the carbon footprint of research [17]. In a domain where the goal is healthy aging, it would be ironic to pursue computational discovery through environmentally harmful practices. System designers should consider model efficiency, data reuse, distillation, caching, and scheduled rather than continuous retraining. They should also weigh the opportunity cost of computation. A smaller, well-validated model may be more useful than a massive model that is difficult to audit. Sustainability is therefore both an engineering and an ethical concern.

Policy implications extend beyond environmental impact. Regulatory agencies may need to classify recommendation systems according to risk. A system that ranks peptides for laboratory research poses different risks from one that recommends supplements to older adults. Policy could require pre-market review for clinical claims, transparency about training data, and post-market surveillance for adverse events. International coordination is important because peptide products and digital health tools cross borders. At the same time, policy must avoid stifling legitimate research. A balanced approach would encourage innovation while mandating evidence transparency, privacy protection, and accountability. The goal is not to prohibit prediction but to ensure that prediction is embedded in a system of scientific and ethical checks.

8. Case Illustrations and Cross-Domain Comparisons

Consider a hypothetical research platform that integrates public peptide sequences, published assay results, and internal omics data from a laboratory studying skin aging. A researcher queries the system for peptides likely to reduce senescence markers in human fibroblasts. The recommender returns several candidates, each with a confidence score, evidence provenance, and uncertainty profile. One candidate is supported by strong sequence similarity to known

antioxidant peptides but has no human cell data. Another is supported by a small animal study but has poor predicted bioavailability. A third has weak direct evidence but a novel mechanism that could be worth exploring. The system's value lies not in choosing one candidate automatically, but in organizing the evidence so that the researcher can make an informed decision. This example shows why recommendation must be coupled with explanation and uncertainty.

Cross-domain comparisons illuminate design choices. In e-commerce, recommender systems can tolerate many errors because the cost of a bad recommendation is low. In drug discovery, errors are more costly, and validation is slow. In clinical decision support, errors can harm patients. In anti-aging consumer applications, errors may be financially and psychologically harmful, especially when they exploit hope or fear. These differences suggest that a single architecture and evaluation framework cannot serve all contexts. Instead, systems should be modular and configurable, with different thresholds for evidence, different explanation styles, and different oversight requirements. The same multimodal backbone could support multiple deployment modes, but only if governance is mode-specific.

Another comparison comes from personalized nutrition. Recommender systems in nutrition often combine dietary preferences, biomarkers, and behavioral data. They face similar challenges of weak evidence, heterogeneous user responses, and commercial influence. Lessons from that domain include the importance of transparency about evidence quality, the need for human oversight by qualified professionals, and the risk of reducing complex biological processes to simplified scores. Anti-aging peptide recommendation should learn from these failures and successes. It should also avoid the temptation to present prediction as certainty. Aging is multifactorial, and no single peptide or recommendation engine can capture its full complexity [18], [19].

9. Limitations and Future Research

Several limitations constrain current multimodal recommender systems for anti-aging peptides. Data scarcity is severe, especially for human outcomes. Most peptide studies use animal models or cell cultures, and translation to humans is uncertain. Multimodal fusion methods often assume that modalities are aligned and complete, which is rarely true in biomedical settings. Evaluation metrics borrowed from recommender systems, such as ranking accuracy, may not reflect scientific validity or clinical utility. Fairness and robustness are difficult to measure when ground truth is contested. Governance frameworks are still developing, and regulatory oversight varies across jurisdictions. These limitations do not make the approach useless, but they demand humility and rigorous validation.

Future research should focus on several directions. First, standardized ontologies and metadata schemas are needed to make peptide evidence comparable across studies. Second, uncertainty-aware multimodal architectures should be developed and tested in real discovery workflows. Third, participatory design should involve researchers, clinicians, regulators, and older adults in defining what constitutes a useful and acceptable recommendation. Fourth, longitudinal studies should examine how recommender systems influence scientific exploration and consumer behavior over time. Fifth, policy research should explore how to regulate health-related recommendation systems without suppressing beneficial innovation. Finally, sustainability research should evaluate the environmental and economic costs of large-scale multimodal training and deployment.

10. Conclusion

Multimodal recommender systems for predicting anti-aging effects of animal-derived bioactive peptides represent a promising but delicate convergence of artificial intelligence, geroscience, and health infrastructure. They can help organize fragmented evidence, prioritize peptide candidates, and personalize recommendations across research and clinical contexts. However, their value depends on far more than predictive performance. They must handle modality imbalance, data provenance, uncertainty, feedback loops, interpretability, robustness, fairness, privacy, and commercial pressure. They must be embedded in governance structures that assign accountability, require evidence transparency, and protect vulnerable users. They must also be sustainable in computational, environmental, and institutional terms.

The central conclusion is that these systems should be designed as accountable scientific instruments rather than opaque optimization engines. A responsible architecture will be modular, provenance-aware, uncertainty-calibrated, and auditable. A responsible deployment will separate exploratory research from medical advice, monitor for disparities, and support human judgment rather than replace it. A responsible policy environment will encourage innovation while enforcing evidence standards and post-market surveillance. If these conditions are met, multimodal recommender systems may become useful tools for navigating the complex landscape of animal-derived bioactive peptides and aging-related health. If they are not met, they risk amplifying hype, narrowing research, and undermining trust in both artificial intelligence and geroscience.

References

1. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
2. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998–6008.
3. Baltrušaitis, T., Ahuja, C., & Morency, L.-P. (2019). Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2), 423–443.
4. Zhang, S., Yao, L., Sun, A., & Tay, Y. (2019). Deep learning based recommender system: A survey and new perspectives. *ACM Computing Surveys*, 52(1), Article 5. <https://doi.org/10.1145/3285029>
5. Covington, P., Adams, J., & Sargin, E. (2016). Deep neural networks for YouTube recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems* (pp. 191–198). ACM. <https://doi.org/10.1145/2959100.2959190>
6. He, X., Liao, L., Zhang, H., Nie, L., Hu, X., & Chua, T.-S. (2017). Neural collaborative filtering. In *Proceedings of the 26th International Conference on World Wide Web* (pp. 173–182). International World Wide Web Conferences Steering Committee. <https://doi.org/10.1145/3038912.3052569>
7. Rendle, S., Freudenthaler, C., Gantner, Z., & Schmidt-Thieme, L. (2012). BPR: Bayesian personalized ranking from implicit feedback. *arXiv*. <https://doi.org/10.48550/arXiv.1205.2618>
8. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*.

9. Wang, X., He, X., Wang, M., Feng, F., & Chua, T.-S. (2019). Neural graph collaborative filtering. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 165–174). ACM.
<https://doi.org/10.1145/3331184.3331267>
10. Liu, T., Song, X., Dai, X., Ma, H., Muhammad, A., He, Z., & Zhao, K. (2025). Effects of animal-derived peptides on aging-related diseases: a narrative review. *Aging Advances*, 2(4), 160-165.
11. Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
12. Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. *Advances in Neural Information Processing Systems*, 28, 2503–2511.
13. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 220–229). ACM.
<https://doi.org/10.1145/3287560.3287596>
14. Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92. <https://doi.org/10.1145/3458723>
15. Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 33–44). ACM.
<https://doi.org/10.1145/3351095.3372873>
16. 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.
<https://doi.org/10.1007/s11023-018-9482-5>
17. Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44–56. <https://doi.org/10.1038/s41591-018-0300-7>
18. López-Otín, C., Blasco, M. A., Partridge, L., Serrano, M., & Kroemer, G. (2013). The hallmarks of aging. *Cell*, 153(6), 1194–1217. <https://doi.org/10.1016/j.cell.2013.05.039>
19. Fontana, L., Partridge, L., & Longo, V. D. (2010). Extending healthy life span—from yeast to humans. *Science*, 328(5976), 321–326. <https://doi.org/10.1126/science.1172539>
20. Choi, E., Bahadori, M. T., Schuetz, A., Stewart, W. F., & Sun, J. (2016). Doctor AI: Predicting clinical events via recurrent neural networks. In *Proceedings of the 1st Machine Learning for Healthcare Conference* (pp. 301–318). PMLR.