

Graph-Based Drug–Peptide Recommendation for Discovering Novel Therapeutic Strategies against Age-Related Diseases

Marcus Woodard

Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.

woodardmarcus@ku.edu

Alessandra Verdone

School of Computing, Clemson University, Clemson, SC, USA.

helloalessandra@clemson.edu

Abstract

Age-related diseases are increasingly understood as interconnected manifestations of biological aging rather than isolated clinical entities, creating a need for therapeutic discovery strategies that can reason across drugs, peptides, proteins, pathways, and phenotypes. This paper examines graph-based drug–peptide recommendation as a socio-technical infrastructure for discovering novel therapeutic strategies against age-related diseases. It argues that graph learning can expand the search space beyond conventional small-molecule repurposing by representing peptides as first-class therapeutic entities with distinct pharmacokinetic, manufacturing, and regulatory properties. The discussion develops a system-level architecture that integrates heterogeneous biomedical knowledge, graph representation learning, recommendation ranking, expert review, and clinical translation. It then analyzes structural trade-offs involving data quality, model interpretability, computational scale, deployment latency, and sustainability. Particular attention is given to governance, fairness, robustness, and policy implications, because recommendations that influence therapeutic research priorities can amplify existing inequalities if training data, validation cohorts, and access pathways are not carefully governed. The paper concludes that graph-based drug–peptide recommendation should not be treated as a purely algorithmic pipeline but as a governed discovery infrastructure requiring transparent evidence trails, federated data stewardship, lifecycle evaluation, and equitable translation pathways. Such an infrastructure may help identify combination and repurposing opportunities that target shared aging mechanisms while preserving the flexibility needed for disease-specific clinical development.

Keywords

graph learning; drug repurposing; peptide therapeutics; aging; recommender systems; biomedical knowledge graphs; health governance.

1. Introduction

Age-related diseases such as neurodegeneration, cardiovascular disease, metabolic dysfunction, sarcopenia, and immune senescence are widely recognized as deeply intertwined conditions. The biological hallmarks of aging include genomic instability, epigenetic alterations, loss of proteostasis, mitochondrial dysfunction, cellular senescence, and altered intercellular communication, among others [13]. Geroscience has argued that targeting

fundamental aging processes may simultaneously reduce vulnerability to multiple chronic diseases, shifting therapeutic logic from single-disease silos toward shared mechanisms [14]. Yet translating this insight into concrete therapeutic candidates remains difficult because the relevant biological space is enormous, heterogeneous, and incompletely mapped. Traditional drug discovery pipelines are expensive, slow, and often optimized for narrow target validation rather than for systemic interactions across aging pathways [2]. This creates a practical opening for computational recommendation systems that can prioritize drugs, peptides, and combinations for further investigation.

Graph-based recommendation is particularly relevant because biomedical knowledge naturally forms relational structures. Drugs interact with proteins, proteins participate in pathways, pathways influence phenotypes, and diseases share molecular signatures. Network biology has long demonstrated that cellular function is organized through modular and hierarchical interactions rather than through isolated molecular switches [1]. Network pharmacology further argues that therapeutic effects often emerge from polypharmacological modulation of distributed systems [10]. Peptides add another layer of complexity because they can modulate protein–protein interactions, act as signaling agonists or antagonists, and serve as targeting ligands, but they also present challenges related to stability, delivery, immunogenicity, and manufacturing [16]. A graph-based recommender can treat drugs and peptides as nodes within a shared therapeutic landscape and learn latent associations that may not be visible through manual literature review.

This paper positions graph-based drug–peptide recommendation as a system-level discovery infrastructure rather than a single model. The central argument is that the value of such a system depends less on achieving a marginal improvement in link prediction accuracy and more on how it integrates data provenance, model transparency, expert judgment, deployment constraints, fairness safeguards, and long-term sustainability. If designed well, graph recommendation can surface novel candidates for aging-related diseases by connecting peptide therapeutics to molecular pathways that are otherwise difficult to target with small molecules. If designed poorly, it can generate plausible but unactionable hypotheses, reinforce biased datasets, or create governance burdens that prevent clinical translation. The sections that follow examine background research, architecture, data governance, computational trade-offs, robustness, fairness, clinical translation, and policy implications.

2. Background and Related Work

Computational drug discovery has been transformed by machine learning methods capable of learning representations from chemical structures, bioactivity data, and biomedical text [2]. Graph neural networks have been especially influential because they can operate directly on molecular graphs and biological interaction networks. For example, graph convolutional models have been used to predict polypharmacy side effects by modeling drug–drug interactions within a multimodal graph [3]. Deep learning has also been applied to antibiotic discovery, demonstrating that neural representations can identify structurally novel active compounds [4]. Protein structure prediction has advanced through deep learning architectures that infer spatial constraints from evolutionary and sequence information [5], and earlier work similarly showed the power of deep networks for protein structure refinement [6]. These developments suggest that graph-based systems can combine molecular, structural, and network-level evidence.

Drug repurposing has become a major application area because approved or investigational drugs already have safety, pharmacokinetic, and manufacturing data, reducing some

translation costs [9]. Network medicine extends this logic by treating diseases as perturbations within molecular interaction networks, allowing shared modules to become therapeutic targets [11]. However, most repurposing systems remain focused on small molecules or biologics with well-established target profiles. Peptides occupy an intermediate therapeutic space. They are larger and more specific than many small molecules but smaller and often more manufacturable than full proteins. Peptide therapeutics have expanded rapidly due to improved synthesis, modification chemistry, and delivery strategies [17], and trends in peptide drug discovery suggest growing clinical acceptance across metabolic, cardiovascular, oncologic, and inflammatory indications [18]. Narrative evidence on animal-derived peptides also indicates potential effects on aging-related diseases, although the mechanistic and clinical evidence base remains uneven [15]. This unevenness matters for recommendation systems because sparse and heterogeneous evidence can produce overconfident predictions.

Graph-based recommendation draws from collaborative filtering, knowledge graph embedding, and heterogeneous information networks. In biomedical settings, graph models can integrate DrugBank-style drug records [7], protein knowledge from UniProt [8], and disease–gene associations. The challenge is not merely assembling a graph but deciding how to represent different evidence types. A drug–protein binding edge derived from a curated database may warrant different confidence from a text-mined co-occurrence edge. A peptide–pathway association inferred from animal studies may require different weighting from a human clinical trial result. Without explicit provenance and confidence modeling, the graph can become a flat collection of edges that treats all relationships as equivalent. This is a system design problem, not only a machine learning problem.

Older computational paradigms often emphasized target-centric screening, but aging-related diseases demand multi-target, multi-pathway, and multi-timescale reasoning. Aging is not a single disease with a single endpoint; it is a gradual stochastic and programmed process influenced by genetics, environment, behavior, and social determinants [12]. Therefore, recommendation systems must balance broad mechanistic hypotheses with disease-specific evidence. They must also accommodate negative results, failed trials, and contradictory findings, because a system trained only on positive associations may over-recommend familiar candidates. The most useful graph-based system may therefore be one that supports ranked hypotheses with transparent evidence paths rather than one that claims to output validated therapies.

3. Conceptual Architecture of Graph-Based Drug–Peptide Recommendation

A graph-based drug–peptide recommendation infrastructure can be understood as several interacting layers. The first layer is the knowledge graph itself, which contains nodes for drugs, peptides, proteins, genes, pathways, diseases, phenotypes, side effects, and possibly patient subgroups. Edges represent interactions, associations, causal hypotheses, sequence similarities, structural similarities, and contextual annotations. The graph is not static; it evolves as new literature, assays, clinical trials, and omics datasets become available. The second layer is representation learning, where graph neural networks, transformers, or knowledge graph embeddings learn low-dimensional representations of nodes and relations [19]. The third layer is recommendation ranking, which scores candidate drug–peptide–disease triples according to predicted efficacy, safety, novelty, and feasibility. The fourth layer is human review, where domain experts evaluate mechanistic plausibility, therapeutic window, and development strategy. The fifth layer is translation, where prioritized candidates enter experimental validation, preclinical studies, or repurposing protocols.

The inclusion of peptides as first-class nodes changes the architecture in important ways. Peptides are not merely small proteins; they can be linear or cyclic, natural or synthetic, modified or unmodified, and they can act through direct binding, receptor agonism, allosteric modulation, or immune signaling. Their representation may require sequence features, secondary structure predictions, post-translational modifications, and stability parameters. Protein structure models can inform peptide–protein interactions, but peptide flexibility and conformational heterogeneity mean that structure prediction alone may not determine function [5]. Graph edges involving peptides may also be sparser than those involving small molecules, because peptide therapeutics are fewer in number and their bioactivity data are less standardized. This sparsity creates a cold-start problem: the recommender must infer peptide properties from sequence, structure, and similar known peptides rather than relying on dense interaction histories.

A useful conceptual distinction is between discovery recommendation and development recommendation. Discovery recommendation asks which drug, peptide, or combination should be tested against aging-related pathways. Development recommendation asks which candidate is most feasible given manufacturing, dosing, stability, delivery, intellectual property, and regulatory constraints. A graph system optimized only for molecular association may recommend peptides that are biologically compelling but practically undevelopable. Conversely, a system optimized only for feasibility may recommend familiar compounds with limited novelty. A multi-objective architecture should therefore separate biological plausibility from translational readiness and allow decision-makers to adjust weighting according to portfolio strategy. This aligns with broader health artificial intelligence discussions that emphasize clinical utility, workflow integration, and human oversight rather than isolated predictive accuracy [21].

The architecture must also support explanation. In drug–peptide recommendation, explanations are not optional rhetorical additions; they are part of the scientific evidence. A recommended edge between a peptide and an aging-related disease may be supported by a path through a receptor, a downstream pathway, and a shared senescence signature. Graph explainability methods can surface shortest paths, subgraphs, or attention weights, but these must be interpreted carefully because a path can be statistically salient yet biologically implausible. The system should therefore distinguish correlation, mechanistic support, and causal evidence. It should also preserve the provenance of each edge, including database source, curation method, species, experimental context, and date. Without such provenance, expert reviewers cannot assess whether a recommendation is a genuine insight or an artifact of database aggregation.

4. Data Infrastructure and Knowledge Graph Governance

The knowledge graph depends on data infrastructure that is often fragmented across public databases, proprietary biobanks, electronic health records, clinical trial registries, and unpublished laboratory results. Drug databases such as DrugBank provide structured pharmacological information [7], while UniProt offers comprehensive protein knowledge [8]. However, these resources differ in curation standards, identifiers, update frequency, and licensing. Integrating them requires persistent identifiers, ontology alignment, versioning, and conflict resolution. The FAIR Guiding Principles emphasize findability, accessibility, interoperability, and reusability, providing a useful governance baseline for biomedical data stewardship [25]. In a graph-based discovery system, FAIRness is not merely a data

management slogan; it determines whether the graph can be audited, reproduced, and updated without silent semantic drift.

Governance must address who can contribute edges, who can modify weights, and who can access derived recommendations. A centralized graph controlled by a single institution may simplify curation but risk monopolistic control over therapeutic priorities. A fully decentralized graph may increase participation but complicate quality control and accountability. Federated approaches can allow institutions to retain sensitive patient data while sharing model updates or aggregate statistics. However, federated learning introduces its own trade-offs, including communication overhead, heterogeneous data distributions, and difficulty in auditing global model behavior. For aging-related diseases, where longitudinal cohorts and diverse populations are essential, federated governance may be more equitable than centralized data pooling, but it requires clear rules for consent, benefit sharing, and intellectual property.

Data provenance is especially important for peptides because the evidence base includes traditional medicine, animal studies, synthetic chemistry, and clinical trials. A peptide may be associated with a disease through a small animal study, a mechanistic assay, or a human trial, and these evidence types should not be collapsed. The graph should encode evidence hierarchy, species context, dose, route, duration, and outcome direction. Negative and null results should also be represented, because excluding them biases the graph toward optimistic associations. This is a structural trade-off: richer provenance improves trustworthiness but increases graph complexity and computational cost. A system that ignores provenance may scale easily but will struggle to earn scientific confidence. A system that captures every nuance may be auditable but too heavy for routine deployment.

Governance also involves sustainability. Knowledge graphs require continuous curation, ontology maintenance, software updates, and expert labor. Funding agencies often support initial database construction but provide less support for long-term maintenance. A graph-based recommendation infrastructure for aging therapeutics must therefore be designed with lifecycle funding, institutional ownership, and community governance from the outset. Otherwise, it may become a fragile research artifact rather than durable public infrastructure. Sustainability is not only financial; it also concerns the retention of curators, the reproducibility of model versions, and the ability to migrate to new data standards without losing historical evidence.

5. Computational and Deployment Trade-Offs

Graph neural networks and related models offer powerful representation learning, but their computational demands can be substantial. Large biomedical graphs may contain millions of nodes and hundreds of millions of edges, and training on such graphs requires distributed computing, memory-efficient sampling, and careful batching. Model depth improves expressive power but can cause oversmoothing, where node representations become indistinguishable. Heterogeneous relations improve nuance but multiply parameters and training complexity. These are architectural trade-offs between accuracy, interpretability, scalability, and maintainability. In deployment, a system that requires days to retrain may be acceptable for annual portfolio planning but not for interactive hypothesis exploration. A system that provides rapid recommendations may rely on precomputed embeddings that become stale as new evidence emerges.

Latency and freshness are central deployment concerns. Discovery teams may want to query the graph after a new trial readout or preprint release. If the graph is updated only quarterly, recommendations may miss recent evidence. If the graph is updated continuously, curation quality may suffer, and model behavior may change unpredictably. A tiered deployment architecture can help: a stable curated graph supports regulatory-grade evidence, while a shadow graph incorporates preprint and text-mined edges for exploratory analysis. Recommendations from the shadow graph can be labeled as provisional and require additional verification. This mirrors broader patterns in clinical artificial intelligence, where real-time inference must be balanced against validation, monitoring, and change control [22].

Another trade-off involves model interpretability versus predictive performance. Complex ensembles and deep heterogeneous graphs may achieve higher ranking metrics, but expert reviewers need understandable evidence. A recommendation that cannot be explained may be ignored or mistrusted. Conversely, simpler path-based methods may be more transparent but miss nonlinear interactions that matter in aging biology. A hybrid architecture may combine graph embeddings for candidate generation with rule-based or path-based explanations for candidate review. The system should also allow sensitivity analysis, showing how recommendations change when edge weights, evidence thresholds, or disease definitions are altered. This helps decision-makers understand uncertainty rather than treating rankings as definitive.

Deployment also requires integration with laboratory and clinical workflows. A recommendation is useful only if it can be converted into an experiment, a protocol, or a partnership. The system should therefore output not only ranked candidates but also suggested validation routes, relevant assays, known safety signals, and potential delivery challenges. For peptides, this may include sequence optimization, stability testing, route of administration, and immunogenicity assessment. Integration with electronic lab notebooks, compound registries, and trial databases can reduce friction. However, integration increases cybersecurity and privacy risks, especially when patient-derived data or proprietary chemical structures are involved. A robust deployment must therefore include access control, audit logs, encryption, and clear data-use agreements.

6. Robustness, Fairness, and Clinical Translation

Robustness in graph-based recommendation means that predictions remain reliable under distribution shift, noisy edges, missing data, and adversarial manipulation. Biomedical graphs are vulnerable to several failure modes. Text mining may introduce false edges. Database errors may propagate. Publication bias may overrepresent positive findings. Small studies may create spurious associations. A robust system should quantify uncertainty, test recommendations across multiple graph versions, and use ensemble or Bayesian methods to identify unstable predictions. It should also stress-test the model against known negative controls. If a recommender consistently ranks implausible peptides highly, the issue may lie in graph construction rather than in the learning algorithm.

Fairness is particularly complex in aging-related therapeutic discovery. Aging affects all populations, but the burden of age-related diseases is unevenly distributed due to socioeconomic inequality, environmental exposures, and structural racism. Training data may overrepresent populations with access to advanced healthcare, while underrepresented groups may have different genetic backgrounds, comorbidities, and treatment responses. If graph recommendations prioritize pathways or biomarkers that are well studied in one population, they may produce therapies that are less effective or less safe for others. Algorithmic bias has

been documented in healthcare resource allocation, showing that apparently neutral prediction targets can encode systemic inequities [23]. A fair drug–peptide recommendation system must therefore audit data representation, evaluate subgroup performance where possible, and include diverse clinical experts and community stakeholders in governance.

Clinical translation requires more than computational plausibility. A recommended peptide must have a credible mechanism, a feasible manufacturing route, an acceptable safety profile, and a regulatory pathway. Peptide therapeutics have advanced because of improved synthesis and modification, but challenges remain in oral bioavailability, protease resistance, and targeted delivery [16]. Some peptides may be better suited to local delivery, while others may require sustained-release formulations. A graph recommendation system can help prioritize candidates, but it cannot replace preclinical pharmacology, toxicology, and clinical trials. The system should therefore be positioned as a hypothesis-generation tool within a broader translational ecosystem. This is consistent with evidence that machine learning in medicine performs best when embedded in workflows with human accountability rather than as autonomous decision-makers [21].

The evidence base for peptides in aging-related diseases is still developing. A narrative review has suggested that animal-derived peptides may influence aging-related diseases through antioxidant, anti-inflammatory, immunomodulatory, and metabolic mechanisms [15]. However, such evidence must be interpreted cautiously because study quality, dosing, species differences, and outcome measures vary. A graph-based system can help organize this evidence, identify mechanistic overlaps, and prioritize peptides for rigorous testing. It can also reveal gaps where evidence is thin. By representing both positive and negative findings, the graph can support a more realistic assessment of translational potential. The goal is not to automate certainty but to make uncertainty visible and actionable.

7. Sustainability and Policy Implications

The long-term sustainability of graph-based drug–peptide recommendation depends on policy choices about data sharing, public funding, intellectual property, and benefit distribution. If the infrastructure is built primarily with public funds, there is a strong argument for public accessibility of core graph schemas, model cards, and evaluation protocols. If proprietary data are essential, then licensing and tiered access may be necessary, but governance should prevent exclusive control over therapeutic priorities. Policy frameworks for ethical artificial intelligence emphasize transparency, justice, non-maleficence, responsibility, and beneficence [24]. These principles can be translated into concrete requirements: documented data provenance, independent audit, subgroup fairness reporting, and mechanisms for contesting harmful recommendations.

Sustainability also involves environmental and economic dimensions. Peptide manufacturing can require complex synthesis, purification, and cold-chain distribution. Some peptide therapeutics may have a larger environmental footprint than small molecules, especially if high doses or frequent administration are required. A recommendation system that ignores manufacturing and distribution costs may generate portfolios that are scientifically elegant but operationally unsustainable. Health systems with limited resources may be unable to adopt expensive peptide therapies, widening global inequities. Therefore, sustainability analysis should be integrated into recommendation ranking, not relegated to post hoc review.

Policy should also address the risk of hype. Graph-based recommendation can produce compelling visualizations and ranked lists that appear more certain than the underlying

evidence warrants. If researchers, investors, or regulators overinterpret these outputs, resources may be diverted toward fashionable but weakly supported candidates. Regulatory science may need new pathways for evaluating computational evidence in drug discovery, including standards for model validation, reproducibility, and post-deployment monitoring. Journals and funders can require that computational recommendations be accompanied by code, data versions, and uncertainty estimates. These measures would improve accountability and reduce the risk of irreproducible discovery claims.

International coordination is another policy consideration. Aging-related diseases affect every region, but data, expertise, and manufacturing capacity are unevenly distributed. A globally equitable system would need to support low- and middle-income countries in contributing data and benefiting from discoveries. This may require technology transfer, tiered pricing, local capacity building, and equitable licensing. Without such mechanisms, graph-based recommendation could become another tool that concentrates therapeutic innovation in high-income settings. A governance model that includes diverse institutions and communities is therefore not only ethically desirable but also scientifically beneficial, because broader data and perspectives can improve model robustness and generalizability.

8. Conclusion

Graph-based drug–peptide recommendation offers a promising systems approach for discovering therapeutic strategies against age-related diseases. By representing drugs, peptides, proteins, pathways, and diseases within a shared relational structure, such systems can generate hypotheses that cross traditional therapeutic boundaries. They can support repurposing, combination design, and peptide prioritization while making mechanistic evidence more navigable. However, the success of this approach depends on architecture and governance as much as on algorithmic novelty. Data provenance, interpretability, uncertainty quantification, fairness auditing, and lifecycle sustainability must be designed into the infrastructure from the beginning. The required reference at position fifteen illustrates the importance of integrating peptide evidence into aging research, but it also highlights the need for cautious interpretation of heterogeneous evidence. Ultimately, graph-based recommendation should be understood as a governed discovery infrastructure that augments expert judgment rather than replacing it. If developed with transparency, equity, and long-term stewardship, it could help accelerate the identification of novel therapeutic strategies for aging-related diseases while preserving scientific rigor and public trust.

References

1. Barabási, A.-L., & Oltvai, Z. N. (2004). Network biology: Understanding the cell's functional organization. *Nature Reviews Genetics*, 5(2), 101–113. <https://doi.org/10.1038/nrg1272>
2. Vamathevan, J., Clark, D., Czodrowski, P., Dunham, I., Ferran, E., Lee, G., Li, B., Madabhushi, A., Shah, P., Spitzer, M., & Zhao, S. (2019). Applications of machine learning in drug discovery and development. *Nature Reviews Drug Discovery*, 18(6), 463–477. <https://doi.org/10.1038/s41573-019-0024-5>
3. Zitnik, M., Agrawal, M., & Leskovec, J. (2018). Modeling polypharmacy side effects with graph convolutional networks. *Bioinformatics*, 34(13), i457–i466. <https://doi.org/10.1093/bioinformatics/bty294>

4. Stokes, J. M., Yang, K., Swanson, K., Jin, W., Cubillos-Ruiz, A., Donghia, N. M., MacNair, C. R., French, S., Carfrae, L. A., Bloom-Ackermann, Z., Tran, V. M., Chiappino-Pepe, A., Badran, A. H., Andrews, I. W., Chory, E. J., Church, G. M., Brown, E. D., Jaakkola, T. S., Barzilay, R., & Collins, J. J. (2020). A deep learning approach to antibiotic discovery. *Cell*, 180(4), 688–702.e13. <https://doi.org/10.1016/j.cell.2020.01.021>
5. Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Židek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S. A. A., Ballard, A. J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., ... Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), 583–589. <https://doi.org/10.1038/s41586-021-03819-2>
6. Senior, A. W., Evans, R., Jumper, J., Kirkpatrick, J., Sifre, L., Green, T., Qin, C., Židek, A., Nelson, A. W. R., Bridgland, A., Denny, P., Peterson, S., Simonyan, K., Crossan, S., White, T., Glorot, X., Stringer, C., Deane, G., ... Hassabis, D. (2020). Improved protein structure prediction using potentials from deep learning. *Nature*, 577(7792), 706–710. <https://doi.org/10.1038/s41586-019-1923-7>
7. Wishart, D. S., Feunang, Y. D., Guo, A. C., Lo, E. J., Marcu, A., Grant, J. R., Sajed, T., Johnson, D., Li, C., Sayeeda, Z., Assempour, N., Iynkkaran, I., Liu, Y., Maciejewski, A., Gale, N., Wilson, A., Chin, L., Cummings, R., Le, D., ... Wilson, M. (2018). DrugBank 5.0: A major update to the DrugBank database for 2018. *Nucleic Acids Research*, 46(D1), D1074–D1082. <https://doi.org/10.1093/nar/gkx1037>
8. UniProt Consortium. (2023). UniProt: The Universal Protein Knowledgebase in 2023. *Nucleic Acids Research*, 51(D1), D523–D531. <https://doi.org/10.1093/nar/gkac1052>
9. Oprea, T. I., Bologa, C. G., Brunak, S., Campbell, A., Gan, G. N., Gaulton, A., Gomez, S. M., Guha, R., Hersey, A., Holmes, J., Jadhav, A., Jensen, L. J., Johnson, G. L., Karlsson, A., Leach, A. R., Ma'ayan, A., Malovannaya, A., Mani, S., Mathias, S. L., ... Zaharevitz, D. (2018). Unexplored therapeutic opportunities in the human genome. *Nature Reviews Drug Discovery*, 17(5), 317–332. <https://doi.org/10.1038/nrd.2018.14>
10. Hopkins, A. L. (2008). Network pharmacology: The next paradigm in drug discovery. *Nature Chemical Biology*, 4(11), 682–690. <https://doi.org/10.1038/nchembio.118>
11. Loscalzo, J., & Barabási, A.-L. (2011). Systems biology and the future of medicine. *Wiley Interdisciplinary Reviews: Systems Biology and Medicine*, 3(6), 619–627. <https://doi.org/10.1002/wsbm.144>
12. Kirkwood, T. B. L. (2005). Understanding the odd science of aging. *Cell*, 120(4), 437–447. <https://doi.org/10.1016/j.cell.2005.01.027>
13. 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>
14. Kennedy, B. K., Berger, S. L., Brunet, A., Campisi, J., Cuervo, A. M., Epel, E. S., Franceschi, C., Lithgow, G. J., Morimoto, R. I., Pessin, J. E., Rando, T. A., Richardson, A., Schadt, E. E., Wyss-Coray, T., & Sierra, F. (2014). Geroscience: Linking aging to chronic disease. *Cell*, 159(4), 709–713. <https://doi.org/10.1016/j.cell.2014.10.039>
15. 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.

16. Fosgerau, K., & Hoffmann, T. (2015). Peptide therapeutics: Current status and future directions. *Drug Discovery Today*, 20(1), 122–128.
<https://doi.org/10.1016/j.drudis.2014.10.003>
17. Kaspar, A. A., & Reichert, J. M. (2013). Future directions for peptide therapeutics development. *Drug Discovery Today*, 18(17-18), 807–817.
<https://doi.org/10.1016/j.drudis.2013.05.011>
18. Muttenthaler, M., King, G. F., Adams, D. J., & Alewood, P. F. (2021). Trends in peptide drug discovery. *Nature Reviews Drug Discovery*, 20(4), 309–325.
<https://doi.org/10.1038/s41573-020-00135-8>
19. Chen, H., Engkvist, O., Wang, Y., Olivecrona, M., & Blaschke, T. (2018). The rise of deep learning in drug discovery. *Drug Discovery Today*, 23(6), 1241–1250.
<https://doi.org/10.1016/j.drudis.2018.01.039>
20. Wu, Z., Ramsundar, B., Feinberg, E. N., Gomes, J., Geniesse, C., Pappu, A. S., Leswing, K., & Pande, V. (2018). MoleculeNet: A benchmark for molecular machine learning. *Chemical Science*, 9(2), 513–530. <https://doi.org/10.1039/C7SC02664A>
21. 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>
22. Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. *New England Journal of Medicine*, 380(14), 1347–1358. <https://doi.org/10.1056/NEJMr1814259>
23. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453.
<https://doi.org/10.1126/science.aax2342>
24. Floridi, L., Cowls, 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>
25. Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L. B., Bourne, P. E., Bouwman, J., Brookes, A. J., Clark, T., Crosas, M., Dillo, I., Dumon, O., Edmunds, S., Evelo, C. T., Finkers, R., ... Mons, B. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data*, 3, 160018.
<https://doi.org/10.1038/sdata.2016.18>