Mechanism-Informed, AI-Driven Frameworks for Discovery and Validation of Aging-Associated Chemical Space

by

Sakshi Arora

(PhD21209)

Under the Supervision of

Dr. Gaurav Ahuja

Department of Computational Biology

Indraprastha Institute of Information Technology, Delhi

New Delhi - 110020

March, 2026

Mechanism-Informed, AI-Driven Frameworks for Discovery and Validation of Aging-Associated Chemical Space

by

Sakshi Arora

A Thesis submitted

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Department of Computational Biology

Indraprastha Institute of Information Technology, Delhi

New Delhi - 110020

March, 2026

Certificate

This is to certify that the thesis titled "**Mechanism-Informed, AI-Driven Frameworks for Discovery and Validation of Aging-Associated Chemical Space**" being submitted by **Ms. Sakshi Arora** to the **Indraprastha Institute of Information Technology-Delhi,** for the award of the degree of **Doctor of Philosophy**, is an original research work carried out by her under my supervision. In my opinion, the thesis has reached the standards fulfilling the requirements of the regulations relating to the degree.

The results contained in this thesis have not been submitted in part or full to any other university or institute for the award of any degree/diploma.

March, 2026

Dr. Gaurav Ahuja

Associate Professor

Department of Computational Biology

IIIT Delhi, 110020

Declaration

I acknowledge that I am fully responsible for the entire content of my thesis, including any sections assisted by online tools, including Artificial Intelligence-based tools. I accept full accountability for any violations of ethical standards in publications arising from the use of such tools.

March, 2026

Sakshi Arora

PhD21209

Dr. Gaurav Ahuja

Department of Computational Biology

IIIT Delhi, 110020

Acknowledgments

First and foremost, I would like to express my deepest gratitude to my late grandparents, Nanuji and Badi Ma, for giving me the strength to embark on my Ph.D. journey. Nanuji's ever-growing scientific temper and Badi Ma's unwavering courage were my strongest sources of inspiration in overcoming the challenges I faced throughout this journey. You both have always exemplified how one should not grow old while aging; my first realization that biological age and how one perceives life and its curveballs matter far more than chronological age. This one is for you both.

I would like to begin by sincerely thanking my Ph.D. supervisor, Dr. Gaurav Ahuja, for his scientific rigor and never-quit attitude. From inspiring me to transition from industry to pursuing a Ph.D., he helped me discover a passion for research that I did not know I possessed. He showed me that even while working in a small lab with limited resources, with strong intent and ambitious dreams, such as conducting impactful research and publishing in high-impact journals, can indeed become a reality. Whether through his inspirational stories or, at times, firm admonitions, he instilled in me a scientific vigor that I know will never age, only mature. His constant support during some truly testing times, along with his patience and understanding, is something for which I will always remain grateful. I also thank him for providing me with opportunities to participate in national and international conferences, which significantly enhanced my scientific knowledge and confidence.

I also express my sincere gratitude to Prof. Ranjan Bose, Director of IIIT-Delhi, for admitting me into the Ph.D. program and granting me access to the institute's resources. I thank IIIT-Delhi for providing high-end computational infrastructure and unwavering institutional support, which enabled me to undertake tasks that would have been otherwise unimaginable. I would like to extend my special thanks to Adarsh Sir and the entire IT team for their prompt and kind assistance whenever I encountered technical difficulties. I also extend my thanks to Ajay Sir, Nidhi Ma'am, and the Purchase Department for facilitating wet-lab procurements and often going out of their way to support us. Their efforts played a crucial role in ensuring smooth progress, especially during revisions and deadlines.

I would also like to acknowledge the administrative staff at IIIT-Delhi, particularly Mrs. Priti Patel, Mrs. Shipra Jain, Mrs. Anshu Dureja, and Mr. Raju Biswas, for always being available to address our queries and resolve academic matters efficiently. I would also like to extend my sincere gratitude to Mr. Imran Khan, Mrs. Sarika, Mr. Mohit, and Mr. Kapil Dev Garg for ensuring the timely disbursement of my fellowship.

I gratefully acknowledge the support of IIIT-Delhi, SERB, and TCS Life Sciences for providing the research fellowship that funded my doctoral studies. I also thank my TCS Mentor, Dr. Kuntal Bhushan, for his kind mentorship.

I would like to thank the Department of Computational Biology and its Head, Dr. G. P. S. Raghava, for their continuous support. I also extend my thanks to all faculty members who generously offered guidance, whether during coursework or while clarifying scientific questions. I am particularly thankful to my internal committee members, Dr. Debarka Sengupta and Dr. Tarini Shankar Ghosh, for their insightful suggestions and motivating words.

I also acknowledge the IIIT-Delhi hostel administrative staff for providing essential amenities. Special thanks to Rajeev Sir, Malti Ma'am, and Shanti Ma'am for their prompt assistance. I am equally grateful to the mess vendor Trilok Uncle, the canteen staff, especially Akshay Bhaiya, and Ravi Tea Stall and its proprietor, Rajesh Uncle, for the much-needed food breaks that sustained me through long days.

I extend my heartfelt gratitude to all my lab members: Aayushi Di, Sanjay, Vishakha Di, Subhadeep, Arushi, Sonam, Suvendu, Shiva, Saveena, and Raidhani, for their constant support, encouragement, and scientific guidance. The days filled with fun, chaos, and collaboration will forever remain close to my heart. I would like to extend my special thanks to Sanjay and Vishakha Di for helping me find my footing in the lab, guiding me through moments of uncertainty, and treating me with immense care and affection.

I am also grateful to my seniors and colleagues: Mansi Di, Asra, Sukriti Di, Gayatri Di, Sagnik, Tharrma, Venky, Manas, Pradeep, Akshaya, Dr. Sumeet Patiyal, Dr. Anjali Dhall, and Dr. Neelam Sharma, whose inspiring Ph.D. journeys made mine feel relatable and whose help was invaluable whenever I needed it. I would also like to thank the M.Tech. and B.Tech. juniors I worked with, especially Aman, Bhavey, and Pragya, who brought enthusiasm and fun to our collaborations while constantly inspiring me with their skills. I also thank Nilesh for helping me brush up on my wet-lab skills and for generously sharing his knowledge.

I am deeply thankful to my entire family, my parents, Mrs. Jyoti Arora and Mr. Sanjeev Arora; my sister, Swati Di; and my jiju, Anuj Arya, for their unwavering faith in me and for being my strongest pillars of support. I would not have been able to complete this journey without you. Being the first in my family to pursue higher studies came with its own set of challenges, yet you never doubted my abilities, for which I will be eternally grateful. I also thank my niece, Mannat, who arrived as a beacon of joy and made this journey far more bearable.

I would like to thank the people closest to my heart, Divya, Mayan, Atish, Mohit, Aanchal, Megha, Nitika, and Niranjan, for holding up a mirror when needed and for their unwavering support. Every Ph.D. journey is unique, as are its struggles, but it was your patience in listening to my rants, along with your love and care, that helped me persevere.

Finally, I would like to thank all the mentors and teachers I encountered during my undergraduate and postgraduate studies who inspired me to learn, grow, and think scientifically. I am also grateful to everyone I met during this journey who, in their own way, shaped my path and inspired me in countless ways.

Sakshi Arora


Abstract

Abstract

Aging is a progressive, multifactorial biological process that drives the risk of nearly all major chronic diseases, including cancer, neurodegeneration, metabolic disorders, frailty, and cardiovascular dysfunction. Although the past two decades have established a molecular framework through the Hallmarks of Aging, translating this knowledge into actionable, small-molecule interventions that enhance healthspan remains a central challenge in the field of geroscience. Experimental discovery pipelines are slow, resource-intensive, and typically explore only a minute fraction of chemical space. Conversely, computational drug discovery approaches, while high-throughput, often rely on chemistry-centric descriptors, exhibit black-box behavior, lack mechanistic interpretability, and rarely generalize to biologically novel molecules. This thesis addresses these long-standing limitations by developing two complementary artificial intelligence systems, AgeXtend and AgeXtend::Mimetics, designed to accelerate mechanism-informed discovery of geroprotective molecules and caloric restriction mimetics (CRMs).

The first objective introduces AgeXtend, a multimodal, bioactivity-driven, and fully explainable AI framework. AgeXtend integrates curated datasets of experimentally validated geroprotectors and neutral compounds with bioactivity-based descriptors, hallmark-specific classification models, toxicity prediction, and target inference modules. By combining mechanistic knowledge with machine learning, AgeXtend achieves robust predictive accuracy across cross-validation, leave-one-out validation, and independent external datasets. Importantly, the explainability module maps predictions onto nine aging pathways, allowing for a mechanistic interpretation of each compound's mode of action. Large-scale screening of ~1.1 billion compounds yielded diverse chemical classes with strong geroprotective potential. Experimental validation confirmed these predictions across three biological systems: *Saccharomyces cerevisiae* chronological lifespan assays, human fibroblast senescence assays, and *C. elegans* lifespan assays. Endogenous metabolites and repurposed drugs predicted by AgeXtend demonstrated lifespan-extending or senomodulatory activity, underscoring the biological fidelity of its predictions.

Building upon this foundation, the second objective presents AgeXtend::Mimetics, a novel computational framework designed to identify Caloric Restriction Mimetics, compounds capable of reproducing CR-like physiological responses without structural similarity to known CRMs. Unlike existing approaches that rely on transcriptomic signatures alone or structural matching, AgeXtend::Mimetics explicitly decouples biological convergence from chemical divergence. Using dual similarity modeling, ridge regression residuals, supervised contrastive learning, and composite CRM fingerprinting, the framework identifies

molecules whose biological signatures align strongly with known CRMs despite having distinct chemical architectures. Large-scale application across thousands of compounds revealed chemically novel, mechanistically plausible CRM candidates that align with pathways such as nutrient sensing, autophagy, mitochondrial remodeling, and metabolic regulation. This framework substantially broadens the chemical landscape of CRM discovery and provides mechanistic clarity on CRM-like effects.

Together, the approaches developed in this thesis demonstrate that explainable, mechanism-oriented AI models can successfully bridge the gap between large-scale chemical exploration and biological relevance. AgeXtend and AgeXtend::Mimetics collectively advance the field of computational geroscience by enabling scalable, interpretable, and experimentally validated discovery of geroprotectors and CRMs. These contributions lay the groundwork for future translational studies, the development of generative design for longevity therapeutics, and the integration of multi-omic datasets to refine mechanism-based discovery pipelines. The thesis highlights both the promise and current limitations of AI in aging biology, providing a roadmap for next-generation computational frameworks that target healthspan extension.

List of Publications (linked)

Thesis Related Publications (2)

1. **Arora S***, Mittal A*, Duari S, Chauhan S, Dixit NK, Mohanty SK, Sharma A, Solanki S, Sharma AK, Gautam V, Gahlot PS, Satija S, Nanshi J, Kapoor N, CB L, Sengupta D, Mehrotra P, Ghosh TS, Ahuja G. Discovering geroprotectors through the explainable artificial intelligence-based platform AgeXtend. *Nat Aging*. 2025 Jan;5(1):144-161. doi: 10.1038/s43587-024-00763-4. Epub 2024 Dec 3. PMID: 39627462. ***(first author)***
2. **Arora S**, Nanshi J, Ahuja G. AgeXtend::Mimetics: A Holistic statistical and AI Framework for discovery of Caloric Restriction Mimetics (CRMs). Manuscript under preparation. ***(first author)***

Other Publications (10; reverse chronological order)

1. Mohanty S. K., Mittal A., Namra, Gaur A., Duari S., Solanki S., Sharma A. K., **Arora S.**, Kumar S., Gautam V., Dixit N. K., Subramanian K., Ghosh T. S., Sengupta D., Gupta S. K., Murugan N. A., Sharma D., Ahuja G. -. Deep Learning Reveals Endogenous Sterols as Allosteric Modulators of the GPCR-Gα Interface. *eLife* 14:RP106397. https://doi.org/10.7554/eLife.106397.1 ***(contributing author)***
2. Kumar S, Solanki S, Gupta M, Mohanty SK, Satija S, Chauhan S, Duari S, Sharma A, Gautam V, **Arora S**, Shome R, Sinha S, Sharma AK, Mittal A, Sengupta D, Murugan NA, Ahuja G. Chemical Dice Integrator (CDI): a scalable framework for multimodal molecular representation learning. *bioRxiv*. 2025 Nov;2025.11.11.687860. doi:10.1101/2025.11.11.687860. ***(contributing author; pre-print)***
3. Gautam V., Duari S., Solanki S., Gupta M., Mittal A., **Arora S.**, Aggarwal A., Sharma A.K., Tyagi S., Pankajbhai R.K., Sharma A., Chauhan S., Satija S., Kumar S., Mohanty S.K., Tayal J., Dixit N.K., Sengupta D., Mehta A., Ahuja G. (2025). scCamAge: A context-aware prediction engine for cellular age, aging-associated bioactivities, and morphometrics. *Cell Reports* 44(2):115270. https://doi.org/10.1016/j.celrep.2025.115270 ***(contributing author)***
4. Solanki S, Mohanty SK, Satija S, Chauhan S, Bandaru NVMR, Dukare S, Tiwari NK, Kumar NR, Aravind AB, Mukherjee S, Gautam V, **Arora S**, Kumar S, Duari S, Sharma A, Shome R, Sengupta D, Abbineni C, Samajdar S, Ahuja G. SynGlue: AI-driven designer for clinically actionable multi-target therapeutics. *bioRxiv*. 2025 Aug;2025.08.28.672835. doi:10.1101/2025.08.28.672835. ***(contributing author; pre-print)***

5. Satija S, Mohanty SK, Jorvekar S, Verma A, Solanki S, Gautam V, **Arora S**, Mittal A, Sharma A, Chauhan S, Duari S, Kumar S, Rai A, Das A, Kakati K, Das K, Sarma A, Tayal J, Sengupta D, Mehta A, Borkar RM, Ahuja G. MutAIverse: an AI-powered, mechanism-backed platform for discovering novel DNA adducts and their precursor genotoxins. *bioRxiv*. 2025 Aug;2025.08.26.672508. doi:10.1101/2025.08.26.672508. ***(contributing author; pre-print)***
6. **Arora, S.***, Satija, S.*, Mittal, A.*, Solanki, S., Mohanty, S. K., Srivastava, V., Sengupta, D., Rout, D., Murugan, N. A., Borkar, R. M., & Ahuja, G. (2023). Unlocking The Mysteries of DNA Adducts with Artificial Intelligence. *ChemBioChem*, 25(1). https://doi.org/10.1002/cbic.202300577 ***(first author)***
7. Mittal, A.*, Mohanty, S. K.*, Gautam, V.*, **Arora, S.***, Saproo, S., Gupta, R., Sivakumar, R., Garg, P., Aggarwal, A., Raghavachary, P., Dixit, N. K., Singh, V. P., Mehta, A., Tayal, J., Naidu, S., Sengupta, D., & Ahuja, G. (2022). Artificial intelligence uncovers carcinogenic human metabolites. *Nature Chemical Biology*, 18(11), 1204–1213. https://doi.org/10.1038/s41589-022-01110-7 ***(co-first author)***
8. Gautam, V., Gupta, R., Gupta, D., Ruhela, A., Mittal, A., Mohanty, S. K., **Arora, S.**, Gupta, R., Saini, C., Sengupta, D., Murugan, N. A., & Ahuja, G. (2022). deepGraphh: AI-driven web service for graph-based quantitative structure–activity relationship analysis. *Briefings in Bioinformatics*, 23(5). https://doi.org/10.1093/bib/bbac288 ***(contributing author)***

Under review

9. Chauhan S*, Satija S*, Solanki S*, Duari S, Gautam V, Rakheja I, Kumar S, **Arora S**, Sharma A, Shome R, Sinha S, Ali SY, Prajapati K, Sengupta D, Murugan NA, Ahuja G. Manuscript under review. 2025. ***(contributing author)***
10. Sharma A, Singh A, Kedia S*, Sharma AK, Gautam V, Sharma P, Khandelwal A, Ramakrishna S, Muddashetty R, Freude K, Solanki S, Chauhan S, Kumar S, Satija S, Duari S, **Arora S**, Gupta A, Shome R, Sengupta D, Nair D, Ahuja G. Manuscript under review. 2025. ***(contributing author)***

Table of Contents

List of Figures

List of Tables

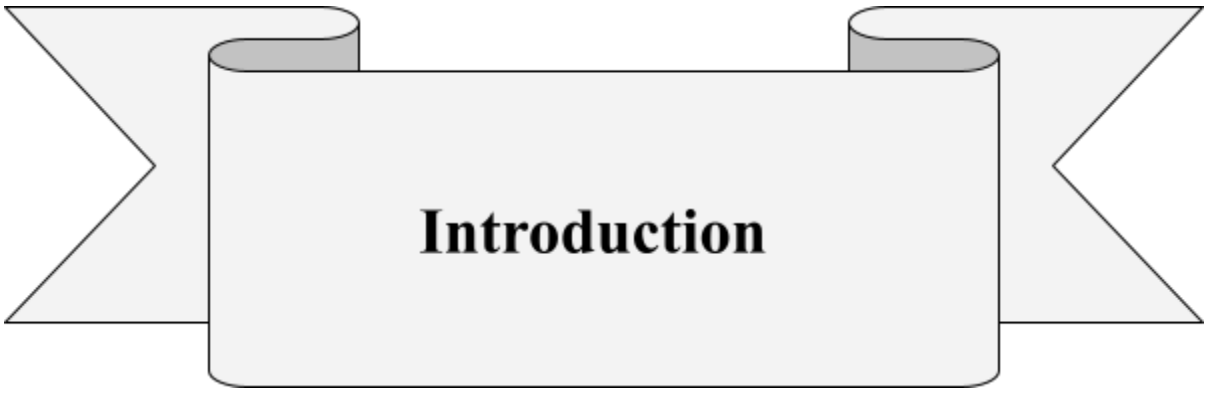
Introduction

Chapter 1 – Introduction

1.1 The Biology and Burden of Aging

Aging is a complex, progressive biological process characterized by the gradual loss of cellular and physiological integrity, resulting in decreased function and increased susceptibility to disease and mortality (López-Otín et al. 2013; Sanada et al. 2025). The processes that lead to this functional decline are not yet fully understood. Across human populations, chronological age is one of the strongest statistical correlates of adult mortality between approximately 30 and 90 years of age, a pattern classically described by the Gompertz relationship (Gompertz 1825; Gavrilov and Gavrilova 2019; Dolejs and Marešová 2017). Yet, despite its robust predictive value at the population scale, chronological age provides limited mechanistic insight into an individual's physiological state or risk of death (Krenz and Strulik 2023). Everyone accumulates chronological years at the same rate, but the pace of biological deterioration varies markedly across individuals. Consequently, two people of the same chronological age can differ substantially in functional capacity, disease burden, and mortality risk. While increases in chronological age are unavoidable, the processes that constitute biological aging are modifiable and reflect a complex interplay between molecular damage, environmental exposures, lifestyle factors, and broader socio-behavioural determinants of health (Ho et al. 2023; Abeliansky et al. 2020).

Biological age provides a quantitative way to capture the inter-individual variability and plasticity of aging (Krenz and Strulik 2023; "Biological Age across the Globe: 1990–2019" 2025). It represents an estimate of an individual's physiological age derived from measurable health or functional indicators, rather than elapsed time. As a personalized and dynamic metric, biological age more accurately reflects the trajectory of health decline than chronological age, making it a stronger indicator for clinical assessment, longevity research, and risk stratification (Z. Li et al. 2023; "Biological Age Predictors" 2017). At the population scale, it also provides a more comprehensive picture of collective health status and can inform evidence-based policy decisions (Levine 2012; Ji et al. 2021).

Aging is not only a biological process but a socioeconomic burden, responsible for the majority of global morbidity and mortality ("Website," n.d.-a; World Health Organization 2015). With the worldwide population aged 60 years and above projected to double by 2050, aging has emerged as the most significant risk factor for chronic diseases, including cancer, cardiovascular disorders, neurodegeneration, and metabolic dysfunction ("Website," n.d.-a, "Ageing and Health," n.d.). In 2015, the World Health Organization declared that the rise in chronic conditions among older adults was a worldwide epidemic (World Health Organization 2015). The presence of multiple comorbidities in the aging population

underestimates the overall burden of age-related diseases. In addition, the pattern of age-related accumulation of burden varies across different demographics, with some populations accumulating age-related damage earlier than others (“Website,” n.d.-a, “Ageing and Health,” n.d., HelpAge International 2024). Substantial differences in aging patterns across different nations also arise due to various factors, such as fertility rates and the speed of change of age-specific mortality, among others(“The Health-Adjusted Dependency Ratio as a New Global Measure of the Burden of Ageing: A Population-Based Study” 2022). This demographic transformation presents unprecedented challenges for healthcare systems, economies, and societal structures worldwide.

While people are living longer, there is little evidence that these additional years are spent in good health (“Website,” n.d.-b, no. 10033, “Ageing and Health,” n.d.). This discrepancy poses a significant burden as it affects individuals' ability to remain active and independent, leading to increased healthcare needs and costs. In low- and middle-income countries, the burden of disease and disability is higher, exacerbating socioeconomic disparities (Huang et al. 2025). Conditions like frailty, which are common in older age, lead to increased vulnerability and dependency, further straining healthcare systems and social support networks (“Website,” n.d.-c). The aging population presents economic challenges, including increased demand for healthcare services, long-term care, and social support systems. This demand places a financial burden on individuals, families, and governments, particularly in countries with limited resources. The cumulative effect of social and economic determinants throughout the life course means that those with the most significant health needs in older age often have the fewest resources to address them, thereby reinforcing existing inequities (“Website,” n.d.-a). Current pharmaceutical interventions primarily address symptoms of age-associated diseases rather than targeting their root molecular causes. This symptom-centric approach results in prolonged treatment durations, reduced efficacy, and increased healthcare expenditures, highlighting the urgent need for interventions that can delay, prevent, or reverse the underlying mechanisms of aging itself (“Website,” n.d.-d).

Therefore, a mechanistic understanding of aging is essential for the development of geroprotective therapies that can extend healthspan rather than merely prolong lifespan, forming the foundation of the present research.

1.2 The Multifactorial Nature of Aging and the Need for Mechanism-Based Interventions

Contemporary aging research has shifted from seeking a single cause, such as a lone gene or system decline, to recognizing aging as an extremely complex, multifactorial process involving the simultaneous interaction of multiple biological, molecular, and physiological mechanisms operating across different

levels of functional organization (Weinert and Timiras 2003). This complex interplay of interconnected cellular and molecular mechanisms operating simultaneously across multiple levels of biological organization is conceived as Hallmarks of Aging (López-Otín et al. 2013). These hallmarks include genomic instability, telomere shortening, mitochondrial dysfunction, epigenetic alterations, deregulated nutrient sensing, cellular senescence, impaired proteostasis, dysbiosis, altered intercellular communication, stem cell exhaustion, and disabled autophagy (López-Otín et al. 2013, 2023). These alterations collectively compromise tissue regeneration, immune competence, and metabolic efficiency, ultimately triggering the onset of various age-associated diseases (Ahmed 2025). These hallmarks are not discrete phenomena; instead, they form a tightly connected network where changes in one domain propagate across others, shaping the tempo and trajectory of organismal aging (Wang et al. 2025).

These cell-intrinsic and cell-autonomous hallmarks converge to disrupt three organizational features of health: spatial compartmentalization (loss of barrier function and tissue integrity), homeostatic maintenance (reduced capacity for sustained metabolic equilibrium), and stress response capability (diminished adaptive and regenerative responses) (“Hallmarks of Health” 2021). Experimentally, the interconnectedness is demonstrated by heterochronic parabiosis, in which simply connecting the circulatory systems of young and old mice triggers systemic rejuvenation of multiple tissues; a phenomenon visualized at the single-cell transcriptomic level, revealing restoration of gene expression across diverse cell types (Pálovics et al. 2022; “Website,” n.d.-e).

These pathways often exhibit compensatory or antagonistic relationships, forming a complex regulatory network that governs cellular health and organismal longevity (López-Otín et al. 2013; Sanada et al. 2025). The multifaceted nature of the aging process thus, necessitates holistic therapeutic strategies rather than conventional single-target drug modalities and also demonstrates that interventions modulating only one pathway are insufficient to meaningfully influence aging trajectories (Wang et al. 2014).

Conventional drug discovery pipelines, which are fundamentally optimized for monogenic diseases through reductionist target-based approaches, fail to adequately address the inherent complexity of aging (Duval 2018). Aging operates through interconnected, pleiotropic mechanisms where interventions produce context-dependent effects that resist capture by traditional single-target pharmacology (López-Otín et al. 2013). This methodological inadequacy is further exacerbated by critical infrastructure deficits in the field: the scarcity of validated geroprotectors to establish proof-of-concept, the absence of validated mechanistic biomarkers to guide drug development, and the lack of standardized experimental assays for aging phenotyping (Herzog et al. 2024; Perri et al. 2025; “Aging: Therapeutics for a Healthy Future” 2020). These foundational gaps have resulted in severely limited chemical diversity among

known geroprotectors, constraining innovation. Furthermore, even promising interventions targeting specific hallmarks; including telomerase gene therapy for telomere attrition, NAD^+ precursors for genomic stability, epigenetic reprogramming with Yamanaka factors, and rapamycin-based autophagy enhancement; encounter substantial translational obstacles: cancer risk from sustained pathway activation, off-target effects from insufficient selectivity, delivery limitations for complex biologics, and critically, insufficient long-term safety data in humans (Bernardes de Jesus et al. 2012; "NAD+ in Aging: Molecular Mechanisms and Translational Implications" 2017; Kane and Sinclair 2019; Szőke et al. 2023; Sanada et al. 2025). The paradoxical challenge of imprecise modulation cannot be overstated: therapies intended to stimulate rejuvenation may inadvertently accelerate aging if not rigorously controlled, triggering cellular senescence, stem cell exhaustion, or malignant transformation. Overcoming these limitations requires fundamentally reimagining the discovery paradigm toward precision, tissue-specific modulation; transient rather than chronic pathway manipulation; and rationally designed, combinatorial approaches capable of achieving multi-hallmark rejuvenation without compromising organismal safety.

Recognition of the multifactorial, interconnected nature of aging has catalyzed the development of geroprotectors: compounds and interventions targeting one or multiple hallmarks of aging to slow, arrest, or reverse the aging process and extend healthspan (Moskalev et al. 2017; "Website," n.d.-f). Geroprotectors broadly aim to address the fundamental causes of aging through diverse mechanisms: genetic intervention (telomerase activation, SIRT6 overexpression), pharmacological inhibition of pro-aging pathways (mTORC1 inhibitors like rapamycin), activation of protective pathways (AMPK activators, sirtuin activators), enhancement of cellular quality control (autophagy inducers), and immune modulation (Gkioni et al. 2025; Chen et al. 2021; Lamming et al. 2013; Mannick and Lamming 2023; Barcena et al. 2025; Zinecker and Simon 2022).

Critically, caloric restriction represents the most robust and widely validated geroprotective intervention across diverse model organisms from yeast to primates, extending lifespan by 10-40% through suppression of nutrient-sensing pathways, activation of autophagy, enhanced mitochondrial function, and reduced inflammation (Taormina and Mirisola 2014; "Dietary Restriction in Rats and Mice: A Meta-Analysis and Review of the Evidence for Genotype-Dependent Effects on Lifespan" 2012). However, the requirement for sustained caloric reduction, maintaining 15-30% reduced intake indefinitely, renders caloric restriction impractical for most humans due to poor compliance (Di Francesco et al. 2024). This practical limitation catalyzed the development of caloric restriction mimetics (CRMs): pharmacological and dietary compounds that replicate the molecular and physiological signatures of caloric restriction without requiring sustained caloric deficit (Dhahbi et al. 2005). Examples include metformin (demonstrating 92% overlap with caloric restriction gene expression profiles), resveratrol

activating sirtuins, rapamycin and its rapalogs inhibiting mTORC1, NAD^+ precursors (nicotinamide mononucleotide, nicotinamide riboside) enhancing mitochondrial function and autophagy, and spermidine inducing autophagy through TFEB pathway activation (Di Francesco et al. 2024; Hofer et al. 2021; Lakshminarasimhan et al. 2013; Lamming et al. 2013). Intermittent fasting and time-restricted eating protocols achieve CRM-like effects while avoiding chronic weight loss (Soliman 2022).

Importantly, CRMs are not entirely distinct from the broader class of geroprotective compounds but instead occupy an overlapping mechanistic space within the aging intervention landscape (Moskalev et al. 2017). Many interventions originally categorized as geroprotectors exert effects that partially recapitulate the signaling programs associated with caloric restriction (Chung and Chung 2019). For example, compounds such as rapamycin, metformin, and spermidine simultaneously influence nutrient sensing, autophagy, mitochondrial metabolism, and cellular stress responses - pathways that represent core nodes within both caloric restriction biology and broader longevity regulation (Mouchiroud et al. 2010). At the systems level, these pathways form an interconnected regulatory network where perturbation of one node propagates across multiple hallmarks of aging (Gabandé-Rodríguez et al. 2019). Consequently, the distinction between caloric restriction mimetics and general geroprotective interventions is often conceptual rather than absolute, with many compounds acting along a mechanistic continuum that modulates shared longevity-associated signaling circuits.

These geroprotectors and CRMs are being systematically evaluated in clinical trials for their capacity to extend healthspan and prevent or reverse age-related diseases, with early evidence in humans supporting improved metabolic parameters, enhanced immune function, reduced chronic inflammation markers, and improved physical and cognitive performance (“Website,” n.d.-g). The convergence of mechanistic understanding of aging hallmarks with therapeutic development of geroprotectors represents a paradigm shift from treating individual age-related diseases toward mechanism-based interventions addressing aging itself as the ultimate risk factor for disease.

1.3 Research Gaps in Gerotherapeutic Discovery

Despite remarkable progress in understanding the biological architecture of aging, the transition from mechanistic insight to effective gerotherapeutic development remains fraught with structural, methodological, and translational barriers. Although numerous studies have demonstrated that modulation of nutrient sensing, mitochondrial homeostasis, proteostasis, inflammation, and other hallmarks can extend lifespan in model organisms, the translation of these interventions to humans has been inconsistent and often limited by safety concerns, context-dependence, and pathway pleiotropy (Kulkarni et al. 2022;

Campisi et al. 2019; Parkhitko et al. 2023; Rolland et al. 2023). Many promising compounds, ranging from CRMs such as metformin and rapamycin to senolytics and NAD^+ boosters, exhibit dose-sensitive or tissue-specific effects that complicate clinical use, reflecting the difficulty of targeting aging-related pathways that are deeply interconnected and essential for organismal homeostasis (Kennedy et al. 2014; Justice et al. 2016; Lee et al. 2024).

A major barrier stems from the reductionist orientation of conventional drug-discovery frameworks, which were designed for monogenic diseases but are poorly suited for a process regulated by dynamic, polygenic, and network-level interactions (Duval 2018). Aging is characterized by tightly coupled hallmarks, such as genomic instability, mitochondrial dysfunction, epigenetic drift, and chronic inflammation, where perturbation of one node often propagates unpredictably across multiple systems (Bell et al. 2019). As a result, single-target interventions can generate compensatory or deleterious off-target effects, limiting their translational feasibility. Systems-level and combinatorial strategies are therefore increasingly recognized as necessary, yet remain underdeveloped (Partridge 2010; Rose et al. 2010).

The discovery pipeline is further constrained by the scarcity of validated geroprotective compounds and the narrow chemical diversity of those already identified ("Dietary Polyphenols as Geroprotective Compounds: From Blue Zones to Hallmarks of Ageing" 2025). Analyses of chemical space reveal that known geroprotectors cluster into a restricted set of scaffolds with overlapping mechanisms, reflecting both biological bias and methodological limitations in screening approaches (Moskalev et al. 2016; Rivero-Segura et al. 2024). Compounding this challenge is the limited reproducibility of many reported geroprotective effects, especially in mammalian systems, where lifespan and healthspan outcomes remain sensitive to strain, diet, environmental variation, and experimental design (Dakik et al. 2019; Furrer and Handschin 2023). These issues underscore the need for systematic, and standardized screening platforms that can interrogate multi-hallmark biology across diverse model systems (Pan et al. 2025).

Another critical gap lies in the absence of robust biomarkers and validated endpoints for evaluating aging interventions in humans (Moqri et al. 2024). While DNA methylation "clocks," transcriptomic signatures, and composite biomarker panels hold promise, many lack mechanistic specificity, sensitivity to interventions, or validation across different tissues and populations (Tao et al. 2024; Perri et al. 2025; Zheng et al. 2024). The absence of reliable, clinically actionable biomarkers complicates trial design, slows regulatory progress, and limits the ability to benchmark different classes of geroprotectors (Rolland et al. 2023; Herzog et al. 2024). Furthermore, aging is not yet recognized as a treatable indication by regulatory agencies, which requires most trials to target individual age-related diseases rather than the

aging process itself ("Advancing Geroscience Research – A Scoping Review of Regulatory Environments for Gerotherapeutics" 2025).

Finally, long-term safety remains a pervasive uncertainty. Interventions that modulate fundamental biological processes, such as telomerase activation, mTOR inhibition, epigenetic reprogramming, or immune rejuvenation, carry inherent risks of oncogenesis, immune dysregulation, metabolic imbalance, or loss of tissue identity (Sanada et al. 2025; Ji et al. 2023; An et al. 2025; X. Li et al. 2023). Understanding these trade-offs, identifying safe therapeutic windows, and designing precision strategies that restrict intervention to specific tissues, dosages, or timeframes represent ongoing challenges in the field (Tenchov et al. 2024).

Collectively, these research gaps underscore the pressing need for next-generation discovery frameworks that integrate systems biology, unbiased chemical space exploration, high-resolution biomarkers, standardized screening platforms, and computational AI-driven prioritization (Groff-Vindman et al. 2025; Marino et al. 2023; Kattuparambil et al. 2025; Lu et al. 2025). Addressing these limitations will be crucial for advancing geroprotective candidates from conceptual promise to clinically viable therapeutics that can extend human healthspan.

1.4 Thesis Objectives

Artificial intelligence offers a powerful opportunity to accelerate the discovery of gerotherapeutics, provided that the computational frameworks are not only predictive but also mechanistically interpretable and biologically grounded. This thesis is therefore driven by the need to create integrated, mechanism-informed AI systems that bridge biological knowledge, chemical diversity, and experimental validation.

Objective 1 focuses on establishing an explainable, bioactivity-centered AI framework, AgeXtend, for the systematic discovery of geroprotective molecules **(Figure 1.1)**. The objective encompasses predictive modeling, hallmark-level mechanistic inference, toxicity profiling, and putative target identification. The underlying hypothesis for Objective 1 is that bioactivity-driven representations encode cross-pathway mechanistic signatures of geroprotection more effectively than conventional chemical descriptors, thereby enabling accurate, interpretable, and biologically meaningful predictions.

Objective 2 aims to develop a statistically principled and biologically convergent–chemically divergent discovery pipeline, AgeXtend::Mimetics, for identifying novel Caloric Restriction Mimetics (CRMs) **(Figure 1.1)**. This pipeline integrates dual cosine similarity, ridge-regression residual calculations, and a

supervised contrastive learning (SupCon) module trained on CRM-associated pathways to construct functional CRM fingerprints. The core hypothesis for Objective 2 is that compounds demonstrating high biological similarity but low chemical similarity to known CRMs exhibit conserved mechanistic attributes of caloric restriction and thus represent strong candidates for novel CRM discovery.

1.5 Thesis Contributions and Roadmap

This thesis makes several conceptual, methodological, and translational contributions to the field of geroscience and AI-driven drug discovery. At a conceptual level, it introduces a unified framework for mechanism-informed, explainable AI applied to aging biology, shifting the paradigm from purely predictive models toward systems that provide biological insight and translational relevance. Methodologically, the thesis presents AgeXtend, a multimodal AI platform that integrates bioactivity-derived features, hallmark-level mechanistic decomposition, toxicity assessment, and target inference into a cohesive pipeline. The work also contributes an innovative statistical and AI-driven CRM discovery framework, AgeXtend::Mimetics, which operationalizes the principle of biological convergence and chemical divergence using dual similarity metrics, ridge regression residuals, and a supervised contrastive learning encoder capable of generating pathway-level CRM fingerprints.

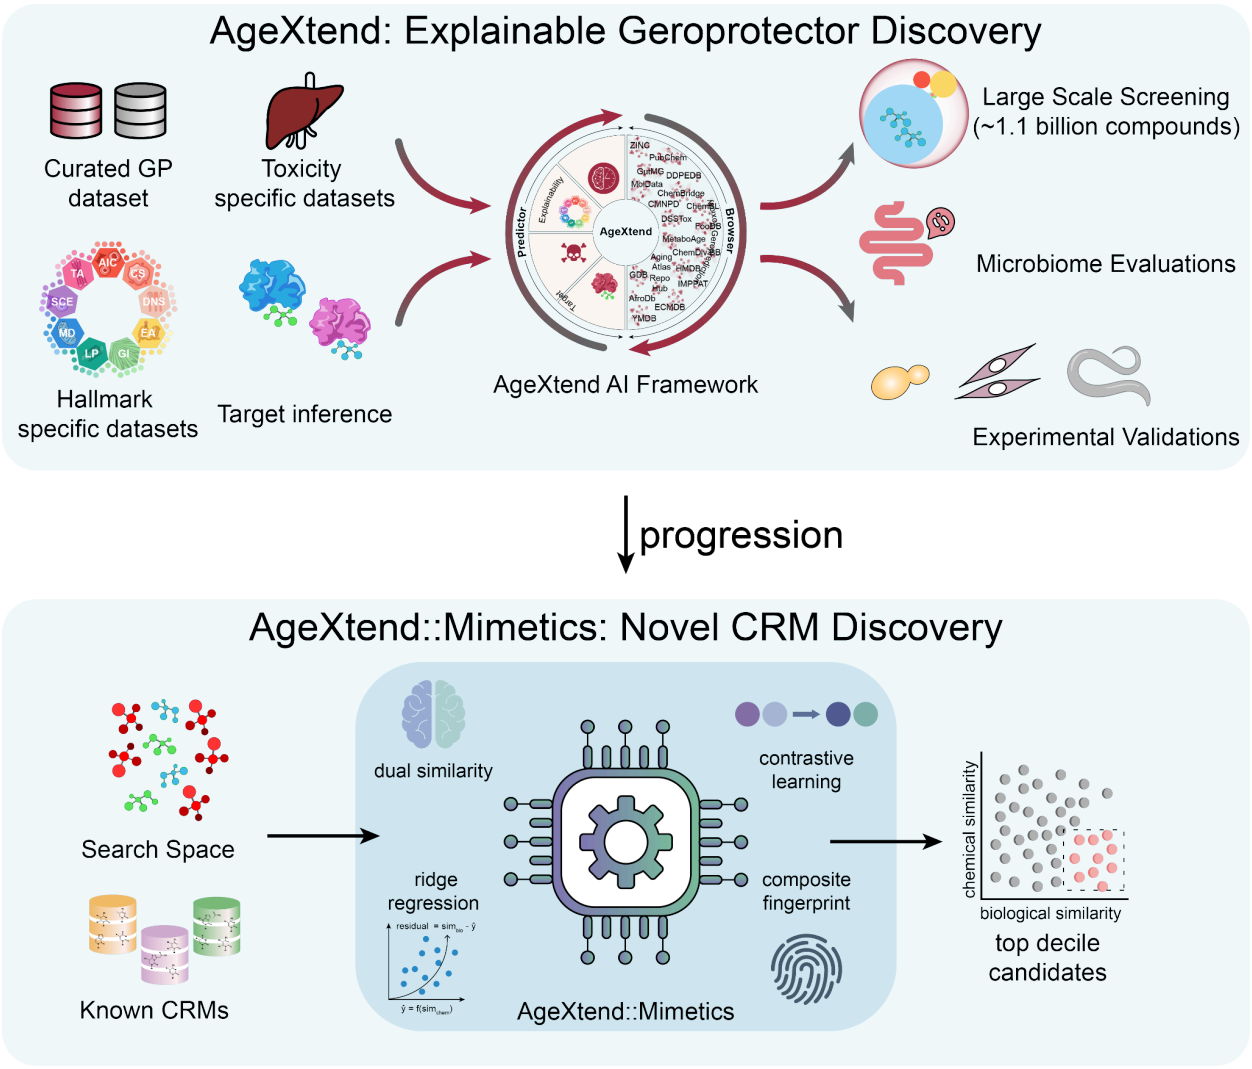

Figure 1.1 Graphical representation depicting the thesis objectives - AgeXtend and AgeXtend::Mimetics

Experimentally, the thesis demonstrates cross-species validation spanning yeast chronological lifespan assays, human fibroblast senescence assays, and *C. elegans* lifespan experiments, establishing the functional relevance of AI-predicted candidates. In addition, this research introduces curated datasets, validated pipelines, and reproducible computational resources that can serve as foundational tools for future aging research.

The remainder of this thesis is organized as follows. Chapter 1 provides the motivation, research gaps, objectives, and broad contributions of the work. Chapter 2 surveys the biological, computational, and pharmacological foundations required to contextualize this research, covering the hallmarks of aging, computational drug discovery methods, microbiome-associated metabolites, and the CRMs landscape. Chapter 3 presents the development, implementation, and validation of AgeXtend. Chapter 4 introduces

the AgeXtend::Mimetics pipeline and details the statistical and AI methods used to identify novel CRMs. Chapter 5 summarizes findings across both objectives and evaluates the broader biological and translational implications. Finally, Chapter 6 concludes the thesis with key insights, limitations, and future directions.

Chapter 2 – Background and Related Work

2.1 Hallmarks of Aging, Metabolic Axes, and Pharmacological Modulation of Aging

Aging is driven by progressive and systemic deterioration in molecular, cellular, and physiological functions. The "Hallmarks of Aging" framework provides a unified conceptual model describing twelve interconnected biological processes: genomic instability, telomere attrition, epigenetic alterations, loss of proteostasis, deregulated nutrient sensing, mitochondrial dysfunction, cellular senescence, stem cell exhaustion, altered intercellular communication, dysbiosis, disabled macroautophagy, and chronic inflammation (López-Otín et al. 2023). These hallmarks interact through dense metabolic and signaling networks, forming conserved longevity axes across species (Mercken et al. 2017; DiLoreto and Murphy 2017; Houtkooper et al. 2010; "Conserved Signaling Pathways Genetically Associated with Longevity across the Species" 2019).

Increasing evidence indicates that the hallmarks of aging are not isolated processes but components of a tightly integrated regulatory network in which perturbations in one domain propagate across others (López-Otín et al. 2013, 2023). For example, mitochondrial metabolism directly shapes epigenetic regulation through the availability of metabolites such as acetyl-CoA, NAD^+, and α-ketoglutarate, which serve as essential cofactors for chromatin-modifying enzymes (Liu et al. 2022). Similarly, defects in DNA repair activate inflammatory signaling through cytosolic DNA sensing pathways such as cGAS–STING, linking genomic instability to chronic inflammation and the phenomenon of inflammaging (Zhao et al. 2023). Nutrient-sensing pathways further integrate metabolic regulation with proteostasis through coordinated control of autophagy and protein quality control systems (Gressler et al. 2023). These examples illustrate that aging hallmarks function as interdependent regulatory modules within a systems-level network rather than discrete biological events.

Central metabolic pathways, including AMPK, mTOR, SIRT1, FOXO transcription factors, and mitochondrial oxidative metabolism, serve as integrative hubs within these hallmarks (Guan et al. 2025). Perturbations in nutrient sensing, mitochondrial function, and proteostasis serve as early drivers of organismal aging, with metabolic axes often dictating tissue-specific outcomes (Amorim et al. 2022; Tenchov et al. 2023). This mechanistic interconnectedness is the basis for pharmacological interventions targeting aging (Guo et al. 2022).

Genomic instability represents a foundational hallmark, reflecting the continuous accumulation of DNA damage from endogenous metabolic byproducts and exogenous insults ("Genomic Instability as a Main Driving Factor of Unsuccessful Ageing: Potential for Translating the Use of Micronuclei into Clinical

Practice" 2021). Failures in DNA repair, chromosomal aberrations, mitochondrial DNA mutations, and defects in nuclear architecture collectively compromise genomic fidelity (Huang and Zhou 2021; Hakem 2008). These alterations not only diminish cellular function but also seed downstream hallmarks, such as senescence and stem cell exhaustion.

Telomere attrition further accelerates genomic vulnerability (Harman and Bryan 2024). Progressive shortening of telomeric DNA impairs chromosome stability, induces replicative senescence, and triggers maladaptive stress responses. Telomere erosion is positioned as a core driver linking replication-associated DNA damage to systemic aging phenotypes, particularly in high-turnover tissues.

Epigenetic alterations constitute another central hallmark, encompassing age-associated changes in DNA methylation, histone modifications, and chromatin organization (Gonzalo 2010; "Epigenetic Interplay between Histone Modifications and DNA Methylation in Gene Silencing" 2008). These epigenetic drifts disrupt transcriptional precision and weaken cellular identity. The crosstalk between epigenetic dysregulation and other hallmarks, especially genomic instability, mitochondrial dysfunction, and loss of proteostasis, is extensive and mechanistically bidirectional ("Mitochondrial-Epigenetic Crosstalk in Environmental Toxicology" 2017).

Proteostasis collapse, characterized by impaired protein folding, clearance, and quality control, is closely linked to the decline of macroautophagy (Taylor and Dillin 2011; "The Aging Proteostasis Decline: From Nematode to Human" 2021). Disabled macroautophagy has a broad role in the turnover of proteins, lipids, and even organelles (Martinez-Lopez et al. 2015). Loss of proteostasis amplifies cellular stress, promotes aggregation-related toxicity, and feeds into mitochondrial damage and inflammatory signaling.

Deregulated nutrient sensing integrates endocrine, metabolic, and stress response pathways (López-Otín et al. 2023). Aging perturbs key nutrient-sensing nodes, IGF-1, mTORC1, AMPK, and sirtuins, leading to imbalanced anabolic and catabolic fluxes. Deregulated nutrient sensing influences multiple other hallmarks, including autophagy, inflammation, mitochondrial health, and DNA repair (Templeman and Murphy 2018). Its centrality is underscored by the number of lifespan-extending interventions that act through nutrient-sensing modulation.

Mitochondrial dysfunction is another pivotal hallmark, reflecting cumulative defects in respiratory chain function, ATP production, mitophagy, and redox balance ("Mechanisms Linking Mitochondrial Dysfunction and Proteostasis Failure" 2020). Mitochondrial damage generates secondary consequences, excessive ROS, apoptotic susceptibility, and metabolic inflexibility, which propagate to genomic

instability, senescence, and inflammation (Peoples et al. 2019). This reciprocal reinforcement renders mitochondrial maintenance a high-value therapeutic target (Reddy and Reddy 2011).

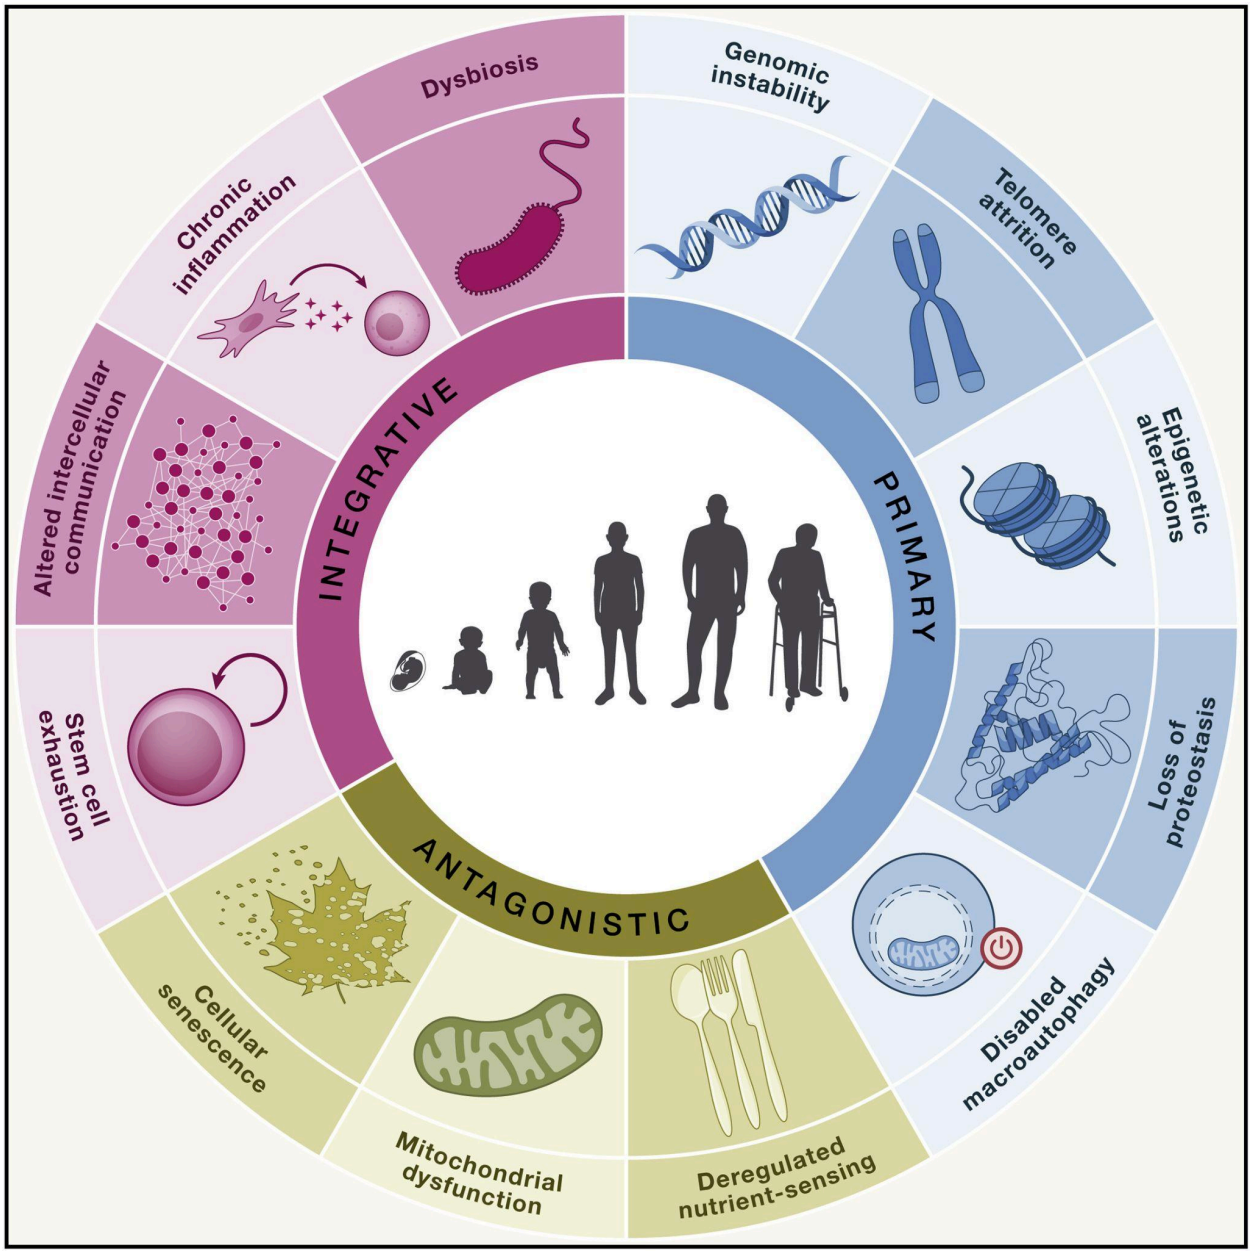

Reproduced from López-Otín et al. 2023 (López-Otín et al. 2023) (permission granted)

Figure 2.1 The hallmarks of aging

This schematic illustrates the twelve proposed hallmarks that collectively define the aging process: genomic instability, telomere shortening, epigenetic drift, impaired proteostasis, reduced macroautophagy, disrupted nutrient-sensing, mitochondrial dysfunction, cellular senescence, depletion of stem cell reserves, altered intercellular signaling, persistent inflammation, and microbial dysbiosis. These hallmarks are organized into three

overarching domains—primary, antagonistic, and integrative—based on their hierarchical roles in the biology of aging.

Cellular senescence emerges from DNA damage, telomere crisis, oxidative stress, and oncogenic stimuli (Nousis et al. 2023). The senescence-associated secretory phenotype (SASP) disseminates pro-aging signals systemically, disrupting intercellular communication and contributing to chronic inflammation, tissue degeneration, and stem cell exhaustion (Itahana et al. 2001; Saretzki 2025). Senescence encompasses both cell-autonomous and non-cell-autonomous effects, consistent with parabiosis and transplanted-cell experiments, which demonstrate the systemic propagation of aging phenotypes (Coppé et al. 2008).

Stem cell exhaustion captures the waning regenerative capacity of tissues, driven by cumulative genomic and epigenetic damage, senescence burden, mitochondrial dysfunction, and altered niche signals (Oh et al. 2014). With age, stem cells lose proliferative vigor and lineage fidelity, contributing directly to functional decline in the hematopoietic, musculoskeletal, and epithelial systems.

Finally, dysbiosis and chronic inflammation form integrative hallmarks that reflect systemic breakdowns in host–microbiome interactions and immune homeostasis (Shen et al. 2025). There is extensive evidence linking age-related microbial shifts to inflammaging, metabolic disruption, impaired microglial function, and cognitive decline (Warren et al. 2024). Chronic inflammation amplifies damage across nearly all other hallmarks, while dysbiosis contributes to immune dysregulation, barrier dysfunction, and altered nutrient processing; exemplifying the multidirectional nature of hallmark interactions (López-Otín et al. 2023).

Several small molecules have demonstrated the ability to modulate these metabolic and hallmark-related processes. Rapamycin inhibits mTOR signaling and promotes autophagy, metformin activates AMPK and enhances mitochondrial efficiency, NAD^+ precursors boost sirtuin activity, and senolytics eliminate senescent cells to reduce inflammaging (Lin et al. 2018; Stephenne et al. 2011; Wang et al. 2021; “Senolytics Decrease Senescent Cells in Humans: Preliminary Report from a Clinical Trial of Dasatinib plus Quercetin in Individuals with Diabetic Kidney Disease” 2019). Autophagy inducers, such as spermidine, promote proteostasis and cytoprotective mechanisms (Yang et al. 2017). Despite these advances, pharmacological strategies remain limited by low mechanistic resolution, modest chemical diversity, and a lack of scalable discovery platforms. This underscores the need for computational systems that can model aging-relevant biological complexity and discover compounds aligned with conserved metabolic and hallmark pathways (Cohen et al. 2022).

2.2 AI in Geroscience and Computational Drug Discovery

AI has emerged as a transformative force in drug discovery, offering mechanistic insights and enabling navigation of vast chemical spaces ("Artificial Intelligence in Drug Discovery and Development" 2021). In geroscience, AI frameworks must capture the multifactorial, pleiotropic, and cross-pathway nature of aging; attributes that challenge conventional prediction models ("Artificial Intelligence (AI) in Drug Design and Discovery: A Comprehensive Review" 2025).

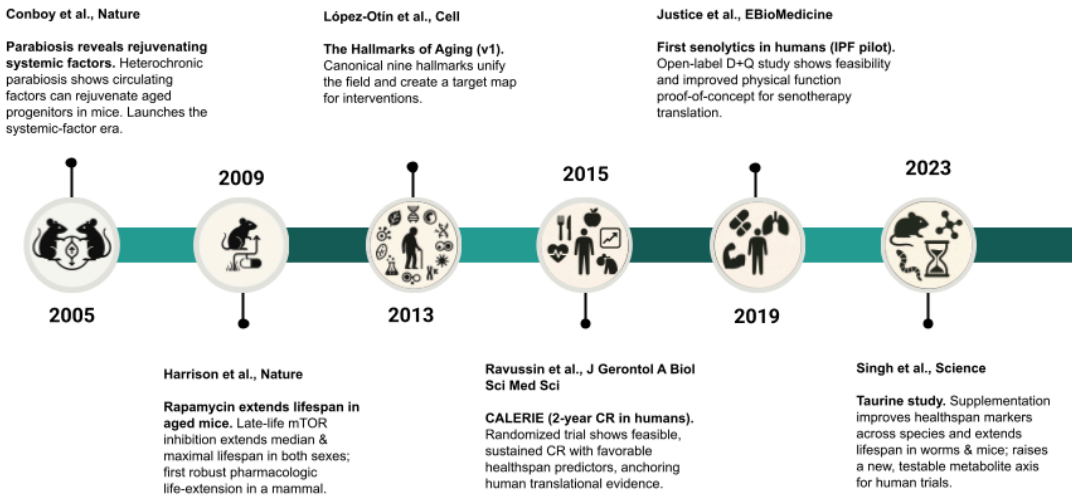

Figure 2.2 Evolution of the modern aging research landscape (2005–2023)

Timeline summarizing key discoveries that transformed aging from a descriptive process into a modifiable biological outcome. Foundational advances include systemic rejuvenation via heterochronic parabiosis (Conboy et al., 2005(Conboy et al. 2005)), late-life lifespan extension by rapamycin (Harrison et al., 2009(Harrison et al. 2009)), and the formalization of the Hallmarks of Aging framework (López-Otín et al., 2013(López-Otín et al. 2013)). Translational validation followed through human caloric restriction (CALERIE, 2015(Ravussin et al. 2015)) and senolytic trials (Justice et al., 2019(Justice et al. 2019)), culminating in recent metabolite-based lifespan modulation such as taurine supplementation (Singh et al., 2023(Singh et al. 2023)).

Several computational paradigms contribute to this space:

A. QSAR and Classical Models

Traditional QSARs utilize handcrafted chemical descriptors and machine learning algorithms, such as logistic regression, ridge regression, random forests, gradient boosting, and SVMs, to predict biological activity. These models offer interpretability and robustness, but are limited when biological complexity exceeds that of chemical structure alone (Santiago-de-la-Cruz et al. 2025; Liu et al. 2024). AgeXtend

incorporates classification and regression methods to predict geroprotective potential, toxicity likelihood, and hallmark alignment.

B. Transcriptomics and Perturbational Signatures

Transcriptomic-based discovery frameworks, including Connectivity Map (CMap) and gene expression reversal methods, match biological perturbation signatures to identify compounds that mimic or reverse specific cellular states (Musa et al. 2018). While powerful, these approaches struggle with cell-line dependency, batch effects, limited mechanistic decomposition, and a lack of integration with chemical representations. AgeXtend::Mimetics extends this paradigm by modeling biological similarity between chemical compounds as a central parameter.

C. Classification & Regression in Geroprotector Prediction

Machine learning (ML) has emerged as a key tool in geroprotector discovery, enabling scalable prioritization of compounds beyond low-throughput experimental screening. Existing computational efforts largely frame the problem as either classification, where compounds are labeled as geroprotective or non-geroprotective, or regression, where quantitative aging-related outcomes are predicted (Moskalev et al. 2016). Both paradigms have been explored in the literature, shaped primarily by data availability and biological complexity.

Early work in geroprotector prediction predominantly adopted classification-based approaches, driven by the scarcity of standardized, continuous effect-size measurements across aging studies (Santiago-de-la-Cruz et al. 2025). Models were typically trained on curated datasets such as DrugAge, using classical supervised algorithms including logistic regression, support vector machines, random forests, and gradient boosting. Most relied on chemical fingerprints or physicochemical descriptors, achieving reasonable performance but often exhibiting limited generalization to chemically novel scaffolds. This limitation reflects both the narrow chemical diversity of known geroprotectors and the inherent uncertainty of negative labels, which frequently correspond to untested rather than truly neutral compounds.

To address these shortcomings, subsequent studies incorporated biologically informed features, such as known drug–target interactions, pathway memberships, and gene expression perturbation profiles (Stolbov et al. 2025; Dönertaş et al. 2018). Transcriptomic similarity to longevity-associated signatures or caloric restriction responses has been particularly influential, improving biological relevance compared to structure-only models. However, such approaches remain constrained by context dependence, incomplete

coverage, and reliance on surrogate cellular systems, which may not fully capture the biology specific to aging.

More recent efforts have explored deep learning architectures, including graph neural networks and pre-trained molecular encoders, to model nonlinear structure-activity relationships (Cai et al. 2022; Yao et al. 2024). While these models can outperform classical methods, their effectiveness in geroscience is limited by small labeled datasets, often necessitating transfer learning or pretraining on large chemical corpora. Interpretability also remains a challenge, as many deep models provide limited mechanistic insight into aging-related predictions.

In contrast, regression-based models aim to predict continuous aging outcomes such as lifespan extension or survival metrics (Mutz et al. 2024). Although regression preserves effect-size information and supports finer prioritization, its application is hindered by the heterogeneity of aging assays, species differences, and experimental conditions. Consequently, regression approaches are often restricted to organism-specific or assay-specific settings or are replaced by ordinal approximations of geroprotective strength.

Overall, existing research highlights persistent challenges in predicting geroprotector, including dataset bias, label noise, and the multifaceted nature of aging **(Table 2.1)**. These limitations have motivated the development of multi-model and mechanism-aware frameworks that explicitly account for biological heterogeneity rather than treating geroprotection as a single homogeneous phenotype. The approaches developed in this thesis build upon these insights by integrating bioactivity-centric representations, interpretable modeling, and mechanistic decomposition, aiming to improve generalization, biological relevance, and experimental utility in geroprotective discovery.

AgeXtend relies on supervised learning and classification to identify geroprotectors and their hallmark level mechanism of action, as well as toxicity and regression to further estimate toxicity risks and dosage. These models integrate bioactivity descriptors (Signaturizer) with chemical features, forming a high-dimensional, biologically enriched representation space (Bertoni et al. 2021). Interpretation using SHAP enhances mechanistic transparency and aligns with the thesis goal of explainable AI.

D. Conditional Mutual Information (CMI) as a Foundation for Mechanism-Aware Similarity Modeling

Conditional Mutual Information (CMI) provides a principled information-theoretic framework for quantifying the dependence between two variables while explicitly controlling for the influence of a third (Thomas M. Cover and Thomas 2006). In biological modeling, this is particularly valuable because

Aspect	Classic ML Models	Deep Learning Models	Bioactivity-Based Models
Typical Algorithms	Logistic regression, SVM, Random Forest, Gradient Boosting, XGBoost	GNNs, CNNs, Transformers, pretrained molecular encoders	Linear/ensemble models or DL on perturbational embeddings
Primary Input Features	Chemical fingerprints (ECFP), physicochemical descriptors (RDKit, Mordred)	Molecular graphs, SMILES, learned embeddings	Drug-induced bioactivity signatures, pathway/target perturbations
Data Requirements	Low–moderate; works with small curated datasets	High; often requires pretraining or transfer learning	Moderate; leverages external perturbation data
Generalization to Novel Chemotypes	Limited; often biased toward known scaffolds	Improved with pretraining, but still chemistry-driven	High; captures biological convergence independent of structure
Interpretability	Moderate (feature importance, SHAP)	Low–moderate (post-hoc explainability)	High; pathway- and mechanism-level interpretability
Handling of Aging Pleiotropy	Weak; treats geroprotection as a single label	Implicit; depends on architecture and data	Strong; explicitly models multi-pathway effects
Typical Use in Geroscience	Early geroprotector classifiers	Emerging, exploratory applications	Mechanism-informed geroprotector/CRM discovery
Key Limitations	Overfitting, label noise, structural bias	Data hunger, black-box behavior	Dependence on quality and coverage of bioactivity data
References	(Santiago-de-la-Cruz et al. 2025), (Yarmolenko and Howlin 2022), (Kapsiani and Howlin 2021)	(Eslami Manoochehri et al. 2018), (Cai et al. 2022)	(Ribeiro et al. 2022), (Yao et al. 2024), (Zeng et al. 2020)

observed associations often arise from confounding factors rather than true mechanistic coupling (Neuberg 2003). By conditioning out known sources of correlation, CMI enables isolation of the residual signal that reflects genuine functional or mechanistic alignment.

Table 2.1 Comparison of Machine Learning Paradigms for Geroprotector Prediction

In this thesis, AgeXtend::Mimetics adopts this philosophy computationally, translating the core intuition of CMI into a tractable similarity-based framework suitable for large-scale chemical screening. The central objective is to disentangle biological convergence from chemical redundancy; that is, to identify molecules that resemble known caloric restriction mimetics (CRMs) in their biological effects, even when they are structurally dissimilar. Formally computing CMI in high-dimensional chemical–biological spaces is often infeasible due to data sparsity and the need for distributional assumptions. Instead,

AgeXtend::Mimetics operationalizes a CMI-inspired surrogate that preserves the conditional reasoning while remaining computationally scalable.

This is achieved through paired similarity measures: cosine similarity in biological embedding space, which captures shared pathway- and mechanism-level signatures, and cosine (or Tanimoto-derived) similarity in chemical space, which captures structural relatedness. Ridge regression is then used to model the expected biological similarity as a function of chemical similarity, effectively learning the baseline level of biological resemblance that can be attributed to shared chemical features alone. The residuals from this regression, defined as the difference between observed and expected biological similarity, represent the component of biological alignment that cannot be explained by chemistry.

These residuals serve as an approximation to conditional mutual information, highlighting compounds whose biological similarity to known CRMs remains high even after conditioning on chemical similarity. In practical terms, this residual-based score prioritizes candidates that are mechanistically aligned yet chemically novel, directly addressing one of the central challenges in CRM discovery. By grounding candidate ranking in this conditional-information framing, AgeXtend::Mimetics avoids over-prioritizing trivial analogs of known CRMs and instead surfaces molecules that expand the explored chemical space while preserving biological relevance.

Overall, this CMI-inspired modeling strategy provides a statistically motivated foundation for the biological convergence/chemical divergence paradigm introduced in this thesis. It underpins the robustness of the CRM prioritization pipeline, supports interpretability at the mechanism level, and ensures that novelty in chemical structure does not come at the cost of biological plausibility.

E. Contrastive Learning and Supervised Contrastive Learning (SupCon)

Contrastive learning is a representation learning paradigm designed to learn meaningful latent spaces by comparing samples relative to one another rather than relying solely on explicit labels. The core principle of contrastive learning is to pull similar samples closer together while pushing dissimilar samples apart in an embedding space, thereby capturing intrinsic structure in the data (Wang and Isola 2020). Originally developed in self-supervised settings, where positive and negative pairs are constructed through data augmentations, contrastive learning has demonstrated strong performance in domains where semantic similarity is more informative than absolute class boundaries (Chen et al. 2020).

In biological and chemical modeling, contrastive learning is particularly appealing because functional similarity often emerges from shared mechanisms rather than identical structures. This makes contrastive objectives well-suited for tasks such as pathway inference, mechanism discovery, and phenotype-driven

clustering, where the geometry of the latent space itself carries interpretive value (Lan et al. 2025; Yang et al. 2024).

Supervised Contrastive Learning (SupCon) extends this paradigm by explicitly incorporating class or mechanism labels to define positive and negative relationships (Khosla et al. 2020). Instead of relying on augmented views of the same sample, SupCon treats all samples sharing the same label as positives and samples with different labels as negatives. This formulation strengthens the alignment between latent representations and known biological categories, producing embeddings that are both discriminative and semantically structured.

In AgeXtend::Mimetics, SupCon is employed to derive CRM pathway fingerprints, embedding mechanistic information about caloric restriction–associated pathways directly into the latent representation. Compounds are labeled according to their known or inferred CRM-related pathway associations, and a neural network encoder is trained to project bioactivity descriptors into a low-dimensional embedding space. A projection head is appended to the encoder, and the supervised contrastive loss is applied to ensure that compounds sharing similar CRM mechanisms are pulled closer together in the embedding space, while mechanistically distinct compounds are pushed apart.

Architecturally, the SupCon framework consists of a backbone network that processes bioactivity-based input features, followed by a projection head that maps intermediate representations into a contrastive embedding space optimized under the SupCon loss. This separation between representation learning and downstream inference allows the learned embeddings to capture mechanistic structure rather than merely optimizing for classification accuracy. After training, the projection head outputs are used to compute pathway-level similarity scores and probabilistic pathway activations, forming the basis of CRM pathway fingerprints.

By structuring the latent space around mechanistic similarity rather than chemical structure, SupCon complements the conditional similarity modeling used elsewhere in AgeXtend::Mimetics. Together, these components ensure that candidate CRMs are prioritized not only for overall biological convergence but also for coherent pathway-level alignment, thereby enhancing interpretability, robustness, and biological plausibility of the CRM discovery pipeline.

Collectively, these AI paradigms shape the computational landscape of modern geroscience, motivating the design choices underlying the two frameworks developed in this thesis.

2.3 Role of Microbiome and Endogenous Metabolites in Aging

The gut microbiome undergoes profound shifts with age, reflecting a transition toward lower diversity, increased pro-inflammatory taxa, disrupted metabolic outputs, and impaired epithelial barrier function (Shen et al. 2025; Warren et al. 2024). Young individuals tend to exhibit metabolite-rich microbiomes, producing short-chain fatty acids (SCFAs), spermidine, indole derivatives, tryptophan metabolites, and specialized lipids that maintain metabolic homeostasis, immune balance, and mitochondrial efficiency (Ghosh et al. 2022a).

Endogenous metabolites have emerged as naturally occurring modulators of longevity pathways (Tomar and Erber 2023). Several microbial metabolites exhibit geroprotective effects by activating AMPK, enhancing autophagy, attenuating mTOR, and the regulation of NAD^+ metabolism (Fang et al. 2023). Cross-species studies demonstrate that microbiome-derived metabolites can extend lifespan in yeast and *C. elegans* while modulating senescence phenotypes in mammalian cells.

AgeXtend's reliance on bioactivity-based descriptors enables recognition of metabolite-like behavioral signatures, allowing cross-domain predictions from endogenous metabolites to synthetic molecules. This positions endogenous metabolites not only as validation benchmarks but also as mechanistically rich anchors for expanding the aging-related chemical space.

2.4 Caloric Restriction Mimetics: Mechanisms, Limitations, and Need for Novel Discovery Frameworks

Caloric restriction (CR) consistently extends lifespan across a wide range of organisms, from yeast to primates. CRMs aim to replicate the systemic benefits of CR: autophagy induction, improved mitochondrial function, enhanced proteostasis, and metabolic rewiring, without imposing dietary constraints (Smith et al. 2010; Mariño et al. 2014; Jafari et al. 2024). Known CRMs such as spermidine, resveratrol, metformin, and NAD^+ boosters activate stress-response pathways, nutrient-sensing arms, and cytoprotective mechanisms ("Recent Advances in the Field of Caloric Restriction Mimetics and Anti-Aging Molecules" 2021; Trisal and Singh 2024).

Despite the broad translational appeal of CRMs, their systematic discovery remains constrained by several scientific, methodological, and conceptual bottlenecks. These limitations collectively explain why, after decades of CR research, only a small set of molecules are recognized as putative CRMs, and even these lack universal validation across organisms and physiological states.

1. Absence of a unified mechanistic definition of CR mimicry

CR influences a constellation of interconnected pathways; autophagy, AMPK–mTOR signaling, mitochondrial biogenesis, redox balance, stress responses, and metabolic flexibility ("Caloric Restriction Mimetics against Age-Associated Disease: Targets, Mechanisms, and Therapeutic Potential" 2019, "Perspectives of the Potential Implications of Polyphenols in Influencing the Interrelationship between Oxi-Inflammatory Stress, Cellular Senescence and Immunosenescence during Aging" 2020). Because CR does not operate via a single hallmark, defining a molecular "CRM signature" has remained elusive. As a result, different laboratories operationalize CR mimicry inconsistently (Flanagan et al. 2020). Moreover, the molecules classified as CRMs often share only a subset of CR phenotypes, not the full systemic profile ("Calorie Restriction Mimetics: Can You Have Your Cake and Eat It, Too?" 2015). This makes the benchmarking of new candidates highly subjective.

2. Limited availability and poor standardization of experimental assays

CRM validation requires multi-layered evidence: transcriptomic responses, metabolomic rewiring, autophagic flux, lifespan extension, and phenotypic rejuvenation across models. Yet, no gold-standard assay exists for systematic CRM screening (Madeo et al. 2014). Most studies rely on organism-specific assays (yeast CLS, worm lifespan, rodent metabolic markers) that are not directly comparable (Lee and Min 2013). Context-dependent variability (diet composition, fasting length, microbiome status) complicates assay reproducibility (Hofer et al. 2021; Ingram et al. 2006). This inconsistency prevents the creation of a harmonized training dataset needed for computational discovery (Ingram et al. 2004).

3. High biological pleiotropy and context dependence

CR triggers tissue- and context-specific adaptations, such as hepatocytes exhibiting metabolic rewiring and increased fatty acid oxidation; neurons emphasizing proteostasis and stress resilience, and muscle prioritizing mitochondrial efficiency ("Caloric Restriction Mimetics against Age-Associated Disease: Targets, Mechanisms, and Therapeutic Potential" 2019). Thus, a molecule may mimic CR in one tissue but fail in another, resulting in false positives or false negatives depending on the assay choice. Moreover, age, sex, feeding schedule, circadian phase, and microbiome state profoundly shape CR responses (Trisal and Singh 2024; Mariño et al. 2014). CRMs often exhibit conditional phenotypes, complicating generalizability across species.

4. Narrow chemical diversity and structural redundancy of known CRMs

Most known CRMs cluster into a few chemotypes: polyamines (spermidine, spermine analogs), phenolic stilbenoids, biguanides, and NAD^+ precursors (Madeo et al. 2014). This limited chemical landscape has

two consequences: discovery efforts repeatedly rediscover close analogs rather than structurally novel compounds, and machine-learning models trained on known CRMs inadvertently learn chemical similarity rather than biological mimicry, decreasing the likelihood of finding new scaffolds. Therefore, a framework that penalizes chemical redundancy is critical.

5. Mechanistic ambiguity and multiplicity of targets

CRMs rarely act through a single target; instead, they modulate broad physiological networks. For example, spermidine induces autophagy while affecting histone acetylation and mitochondrial translation; metformin modulates AMPK, mitochondrial complex I, redox state, and gut microbiota; resveratrol engages sirtuins but also alters inflammatory signaling and ROS metabolism (Madeo et al. 2014; Singh 2025). This mechanistic multiplicity complicates the causal inference, cross-species extrapolation, and predictive modeling, since high-dimensional biological descriptors dilute signals with off-target noise.

6. Sparse, noisy, and fragmented biological datasets

CRM-relevant signatures (transcriptomes, proteomes, metabolomes, phenotypes) are generated under heterogeneous laboratory conditions and are therefore highly batch-sensitive, often lacking matched controls, and are distributed across independent studies with differing endpoints (Bibyk et al. 2021; Lushchak and Gospodaryov 2017). This fragmentation limits the construction of harmonized similarity models and complicates efforts to infer conserved CR-like biological patterns.

7. Lack of frameworks integrating biological and chemical domains

Most discovery pipelines operate within a single domain; chemoinformatics-only screens tend to bias toward structural similarity, while transcriptomics-only screens capture biological relevance but struggle to resolve novelty or drug-likeness. This necessitates the development of a framework that allows mechanistic-driven entanglement with known chemistry.

8. Limited emphasis on mechanistic interpretability

Current CRM screens focus on phenotypic endpoints (e.g., lifespan extension) without elucidating the pathway-level mechanisms (Eggert 2013). Without mechanistic fingerprints, potential candidates are hard to prioritize for translation, while their off-target effects remain unexplored. Clinically relevant CR axes (autophagy, mitochondrial hormesis, NAD^+ metabolism, proteostasis) cannot be systematically mapped.

9. Absence of quantitative scoring systems for "degree of CR mimicry"

CRM classification is currently binary, either CRM or not, rather than graded; however, biological mimicry exists along a continuum (Campisi et al. 2019). Without a quantitative metric, mild mimetics may be overlooked, and false positives may inflate the literature, further complicating the cross-compound comparison.

These barriers necessitate a mechanism-informed, statistically principled framework that captures biological similarity independent of chemical structure, exactly what AgeXtend::Mimetics operationalizes through dual cosine similarity, ridge regression residuals, pathway-level SupCon embeddings, and CMI-driven mechanistic clarity.

2.5 Positioning AgeXtend & AgeXtend::Mimetics in the Global Landscape

Most existing computational strategies for aging interventions either rely on structural similarity, which biases discovery toward chemical analogs, or use prediction-only methods that provide limited mechanistic insight. A small fraction incorporates transcriptomic signatures, but few combine chemical, biological, mechanistic, and statistical reasoning in a unified AI system.

AgeXtend advances the field by integrating predictive modeling, interpretability (hallmark-level explanations), toxicity and target inference, and multi-species validation. Its bioactivity-centric design enables mechanistically grounded predictions that extend beyond structural novelty.

AgeXtend::Mimetics represents a significant conceptual leap by operationalizing the principle of biological convergence with chemical divergence. Through dual-similarity modeling, residual-based ranking, pathway fingerprinting via supervised contrastive learning, and CMI-guided mechanistic evaluation, it provides the first scalable, biologically informed, and structure-independent approach for CRM discovery.

Together, these frameworks reorient AI-driven aging research toward mechanistic transparency, statistical rigor, and translational robustness; addressing longstanding barriers in gerotherapeutic discovery and establishing a new blueprint for the field.

Chapter 3 – Objective 1: Discovering geroprotectors through the explainable artificial intelligence-based platform AgeXtend

3.1 Introduction and Rationale

The discovery of geroprotective molecules requires computational frameworks that can integrate chemical structure, biological activity, mechanistic pathway information, and experimental validation. Traditional drug-discovery pipelines fall short in this domain because they assume monogenic disease mechanisms and rely heavily on linear or reductionist models. However, aging is a fundamentally pleiotropic process driven by the dysregulation of multiple interconnected hallmarks, each influenced by complex metabolic, genetic, and signaling networks (Vega Magdaleno and de Magalhaes 2024). This necessitates discovery systems capable of modeling multidimensional biological interactions.

While artificial intelligence has transformed traditional drug discovery, most AI systems remain black-box predictors optimized for high accuracy without mechanistic transparency (Ferreira and Carneiro 2025). Such models frequently provide strong numerical predictions but lack interpretability, making them poorly suited for biology-centric applications where understanding why a prediction is made is as important as the prediction itself (ŞAHiN et al. 2024). Biological systems demand mechanistic reasoning, particularly in geroscience, where pathway interactions, pleiotropy, and cross-species conservation determine translational success ("Artificial Intelligence for Aging and Longevity Research: Recent Advances and Perspectives" 2019). Consequently, black-box AI models risk generating predictions that are uninterpretable, non-actionable, or biologically irrelevant.

AgeXtend bridges this critical gap by providing a mechanism-informed, explainable AI framework that is explicitly designed for aging biology. It integrates bioactivity-derived descriptors with classical machine-learning models, enabling both high predictive performance and biological interpretability. The framework includes hallmark-level mechanistic mapping, SHAP-based feature attribution, toxicity profiling, and target inference modules, allowing users to connect predictions with aging pathways and translational outcomes. This explainability is essential for prioritizing candidates for downstream experimental testing and ensuring biological plausibility across species.

AgeXtend, therefore, represents a hybrid computational–biological framework that not only predicts geroprotectors but also contextualizes them within the mechanistic framework of aging biology. This chapter presents the dataset construction, feature engineering, model development, interpretability modules, and experimental validation that together constitute the AgeXtend framework.

3.2 Methods

3.2.1 Data Curation and Preprocessing

The AgeXtend framework relies on carefully curated, high-confidence datasets covering geroprotective activity, hallmark-level biological processes, and toxicity endpoints **(Figure 3.1a,b; Extended Data Figure 1a)**. Data compilation was performed manually using authenticated literature and database resources, with extensive cross-verification across independent sources. Strict filtering criteria were applied to ensure mechanistic reliability and reduce noise during model development.

3.2.1 (1) Geroprotective Prediction Dataset

To build the primary geroprotective prediction module, we compiled a binary classification dataset representing compounds with experimentally validated lifespan-modulating effects. The ground-truth dataset consists of 972 unique compounds, including 583 geroprotectors (class 1) and 389 neutral compounds (class 0) **(Table 3.1; Figure 3.1b)**. Geroprotectors (GPs) were curated from: DrugAge (build 4; 20 November 2021 release (Barardo et al. 2017)) and Geroprotective Compounds Database (Barardo et al. 2017; Moskalev et al. 2017) **(Extended Data Figure 1a)**. Both resources enforce stringent inclusion criteria based on reproducibility, assay quality, and organismal evidence. Neutral compounds (N-class) were obtained from peer-reviewed experimental studies reporting no significant effects on lifespan or aging-related phenotypes (Ye et al. 2014; Kornfeld and Evason 2006; Hunt et al. 2011; Evason et al. 2005).

Data curation included removal of: inorganic compounds, salts, solvent components, and metal complexes, duplicates (SMILES-level filtering), and conflicting entries across studies.

To enable robust model evaluation, an independent validation set was assembled by screening ~60 recent publications. This external dataset comprises 74 additional GPs and 10 neutral compounds, serving as a stringent held-out benchmark for assessing generalization.

3.2.1 (2) Hallmark-Level Explainability Datasets

To enable mechanistic inference, we compiled nine hallmark-specific datasets, each consisting of two classes representing positive and negative modulators of a specific biological process **(Table 3.1; Figure 3.2a; Extended Data Figure 1a)**. These datasets correspond to well-characterized pro-aging or anti-aging pathways, including: Genomic instability (GI), Telomere attrition (TA), Epigenetic alterations (EA), Loss

of proteostasis (LP), Deregulated nutrient sensing (DNS), Mitochondrial dysfunction (MD), Cellular senescence (CS), Stem cell exhaustion (SCE), and Altered intercellular communication (AIC).

Class 1 compounds exhibit experimentally validated inhibitory, suppressive, or compensatory actions on the above processes; for example, sirtuin activators (EA, AIC), NF-κB inhibitors (AIC), farnesyltransferase inhibitors (GI), and stem-cell proliferators (SCE) **(Figure 3.2e)**. Class 0 compounds either lack activity or exhibit pro-aging effects within the associated pathway. Conflicts or discrepancies within or between datasets were systematically eliminated.

This hallmark-level curation enables AgeXtend to learn mechanistic patterns directly aligned with aging biology rather than relying solely on bioactivity features for geroprotection prediction.

3.2.1 (3) Toxicity Module Datasets

Long-term safety is essential for geroprotective interventions. To model toxicity and ADMET properties, we constructed 15 dataset modules sourced from vNN-ADMET, spanning key pharmacokinetic and toxicological endpoints: AMES mutagenicity, Mitochondrial membrane potential (MMP) disruption, CYP450 inhibition: CYP1A2, CYP2C19, CYP2C9, CYP2D6, CYP3A4, Hepatotoxicity, HLM stability, hERG channel blockade, Drug-induced liver injury (DILI), Blood–brain barrier permeability (BBB), P-glycoprotein inhibitors/substrates, and Maximum Recommended Therapeutic Dose (MRTD; regression) (Schyman et al. 2017) **(Table 3.1; Figure 3.2g; Extended Data Figure 1a)**.

All classification datasets were binarized into positive (active) and negative (inactive) classes, whereas MRTD was treated as a continuous regression task. These modules allow AgeXtend to perform early-stage safety filtering and prioritization of low-toxicity geroprotective candidates.

3.2.1 (4) Functional Group Analysis

To gain structural insights across all datasets, we conducted a functional group analysis using the ChemmineR library (v.3.46.0). Compounds were converted to SDFset format, and functional group counts were extracted to produce a scaled matrix representing the chemical space across all models. Heatmaps generated from this analysis highlight functional group enrichment patterns and structural motifs characteristic of positive vs. negative biological classes.

This analysis informed interpretability and ensured that model predictions were not driven by trivial chemical motifs or dataset biases.

Table 3.1 Counts of all the Datasets

Geroprediction Module Dataset			
Property	**Count**		
Geroprotectors	583		
Neutral	389		
Explainability Module Datasets			
Property	**Count : 1**	**Count : 0**	**Total**
Altered Intercellular Communication	452	328	780
Cellular Senescence	1292	1257	2549
Deregulated Nutrient Sensing	328	42971	43299
Epigenetic Alteration	422	41372	41794
Genomic Instability	400	5745	6145
Loss of Proteostasis	2184	986	3170
Mitochondrial Dysfunction	871	172	1043
Stem Cell Exhaustion	669	969	1638
Telomere Attrition	209	136	345
Toxicity Module Datasets			
Property	**Count : 1**	**Count : 0**	**Total**
AMES	3502	3008	6510
MMP	1044	5094	6138
CYP1A2	1332	6160	7492
CYP2C19	1516	6567	8083
CYP2C9	938	7060	7998
CYP2D6	1246	6491	7737
CYP3A4	2393	7776	10169
Hepatotoxicity	1968	4113	6081
HLM	1162	2043	3205
hERG	277	385	662
DILI	768	628	1396
BBB	197	156	353
PGP-Inh	1343	961	2304
PGP-Subs	421	399	820
MRTD			1184

3.2.2 Feature Engineering

Accurate prediction requires robust molecular representations that capture both chemical and biological properties. We implemented a multi-tier descriptor extraction strategy across all datasets.

3.2.2 (1) SMILES Standardization

All compounds were represented as Canonical SMILES, generated using OpenBabel (v.3.0.0). Dataset-wide filtering eliminated redundancy and inconsistencies arising from: stereoisomerism, salt forms, protonation inconsistencies, and tautomeric variations. Ensuring one canonical representation per compound was essential for downstream feature extraction and cross-dataset compatibility.

3.2.2 (2) Descriptor Extraction

We benchmarked several descriptor families for the Geroprediction model:

A. Chemistry-Based Descriptors

- Mordred descriptors (v.1.2.0): 1600+ chemophysical, topological, and 2D structural metrics(Moriwaki et al. 2018).
- ECFP fingerprints (1024–2048 bits): circular substructure-based hashing of molecular fragments(Axen et al. 2017).

These descriptors have historically been used in QSAR modeling but often fail to capture biological context relevant to geroprotective activity.

B. Graph-Based Descriptors

- ConvMol (DeepChem v.2.6.2): message-passing graph neural network embeddings encoding local and global structural patterns(GitHub, n.d.).

These descriptors provide structural expressiveness but are insufficient to model biological pleiotropy.

C. Bioactivity-Based Descriptors

- Signaturizer (v.1.1.11): embeddings derived from large-scale perturbational biology datasets, capturing transcriptional and functional similarities between molecules(Bertoni et al. 2021).

3.2.2 (3) Descriptor Performance and Selection

A systematic benchmarking pipeline was conducted to assess descriptor performance across GP prediction. Bioactivity-based descriptors (Signaturizer) consistently outperformed all structural and graph-based descriptors in terms of accuracy, F1-score, ROC-AUC, and cross-validation robustness. Signaturizer embeddings also improved hallmark interpretability because they encode biological signals rather than solely structural or physicochemical information. Consequently, Signaturizer was selected as the primary descriptor space for all downstream AgeXtend models.

3.2.2 (4) Handling Missing Values

Missing descriptor values were imputed using feature-wise mean substitution, performed after dataset standardization. This ensured consistent input dimensionality across all datasets and prevented downstream learning biases.

3.2.3 Geroprediction Model Development and Training Pipeline

The AgeXtend framework required building multiple predictive models spanning geroprotective activity, hallmark-level biological processes, toxicity endpoints, and target inference. To ensure performance consistency and methodological rigor across all modules, we implemented a standardized model generation pipeline comprising (i) data balancing, (ii) classifier benchmarking, (iii) hyperparameter optimization, (iv) feature selection, and (v) cross-validation stability checks.

3.2.3 (1) Class Balancing and Preprocessing

Several datasets within AgeXtend exhibited class imbalance, particularly in hallmark-specific and toxicity-related bioactivity profiles. To mitigate performance distortions arising from skewed class distributions, we employ the Synthetic Minority Oversampling Technique (SMOTE) (Blagus and Lusa 2013). SMOTE was chosen for its ability to synthesize plausible minority-class feature vectors while preserving the underlying data manifold. Oversampling was performed with a proportion of 0.8 for optimal balance between synthetic diversity and overfitteding avoidance.

3.2.3 (2) Classifier Benchmarking and Selection

A broad panel of classical and nonlinear machine-learning classifiers was systematically evaluated for each dataset: Logistic Regression (LR), Support Vector Machines (SVMs), Random Forest (RF), Extra Trees (ET), Gradient Boosting Classifier (GBC), Gaussian Naive Bayes (GNB), k-Nearest Neighbors (kNN), Stochastic Gradient Descent (SGD), XGBoost (XGB), and Multi-layer Perceptron (MLP).

For each model, performance was assessed using fivefold cross-validation, with accuracy, ROC–AUC, Matthews correlation coefficient (MCC), F1 score, precision, recall, and Cohen's κ serving as evaluation metrics. The classifier achieving the highest performance across folds was chosen as the base learner for that specific dataset.

3.2.3 (3) Hyperparameter Optimization and Feature Selection

To refine model performance, we conducted randomized grid-search hyperparameter optimization, exploring a diverse parameter space for each classifier. The best hyperparameter configuration was subsequently evaluated for stability via tenfold cross-validation, ensuring that performance gains were not artifacts of particular training splits.

Feature selection was integrated into the optimization workflow using the Boruta algorithm(Kursa and Rudnicki 2010), which identifies all relevant features contributing to prediction by comparing feature importance against randomized shadow features. Boruta enabled extraction of the most informative molecular descriptors, reducing dimensionality while improving both interpretability and predictive robustness.

3.2.3 (4) Final Model Training, Evaluation, and Cut-off Determination

Once the optimal classifier, hyperparameters, and feature subset were determined, the final model was trained using the complete dataset. Leave-one-out cross-validation (LOOCV) was additionally performed for the Geroprediction model to quantify model sensitivity to individual data points.

For the geroprotective prediction module, we computed a stringent decision threshold for predicted probabilities. This cut-off was determined by evaluating false-positive rates across thresholds from 0.50 to 0.99 in each fold of the tenfold cross-validation. The optimal cut-off was defined as the mean threshold achieving zero false positives across all folds, ensuring high-confidence predictions for downstream analyses.

3.2.3 (5) Descriptor Benchmarking for Geroprediction Module

To benchmark descriptor families for the Geroprediction module, we trained models using: Mordred descriptors, ECFP fingerprints (1024 and 2048 bits), Graph-based descriptors (ConvMol, PMG, MolGraphConv), and Bioactivity-based descriptors (Signaturizer). Identical train–test splits were used to eliminate sampling bias. No feature selection was applied during descriptor benchmarking to ensure fair comparison.

Visual inspection using t-SNE projections (n_components=2, perplexity=30) revealed that bioactivity descriptors produced the clearest separation between experimentally annotated GPs and N compounds. Consistently, Signaturizer outperformed all other descriptor families across every metric, validating its selection as the primary representation for all of AgeXtend's modules **(Table 3.2).**

Table 3.2 Model parameters for Geroprediction and Explainability Modules

Module	Property	ML Model	Class Balancing	Model Parameters
Geroprediction	Geroprediction	Support Vector Machine	Up : 0.8	C=1.5, gamma=2.5, probability=True
Explainability	Altered Intercellular Communication	Support Vector Machine	Up : 0.5	C=2, gamma=1, probability=True
Explainability	Cellular Senescence	Support Vector Machine	Up : 0.5	C=1.5, gamma=2.5, probability=True
Explainability	Deregulated Nutrient Sensing	Support Vector Machine	Down	C=1.5, gamma=1, probability=True
Explainability	Epigenetic Alteration	Support Vector Machine	Up : 0.5	C=1.5, gamma=2.5, probability=True
Explainability	Genomic Instability	Random Forest	Up : 0.8	bootstrap=False, max_depth=64, min_samples_split=4, n_estimators=108
Explainability	Loss of Proteostasis	Support Vector Machine	Up : 0.8	C=2, gamma=2.5, probability=True
Explainability	Mitochondrial Dysfunction	XGBoost	Up : 0.8	base_score=0.5, booster='gbtree', colsample_bylevel=1,colsample_bynode=1, colsample_bytree=0.8,enable_categorical=False, gamma=0.1, gpu_id=-1,importance_type=None, interaction_constraints='',learning_rate=0.2, max_delta_step=0, max_depth=6,min_child_weight=2, missing=nan, monotone_constraints='()',n_estimators=97, n_jobs=20, num_parallel_tree=1, predictor='auto',random_state=0, reg_alpha=0, reg_lambda=1, scale_pos_weight=1,subsample=0.5, tree_method='exact', validate_parameters=1,verbosity=No

				ne
Explainability	Stem Cell Exhaustion	Extra Tree Classifier	Up : 0.8	n_estimators=5000
Explainability	Telomere Attrition	Extra Tree Classifier	Up : 0.5	n_estimators=500

3.2.4 Explainability, Toxicity, and Target Inference Modules

Beyond accurate prediction, AgeXtend was explicitly designed to provide mechanistic transparency and translational relevance. The explainability module, toxicity module, and target inference module form the interpretive backbone of the framework.

3.2.4 (1) Explainability Module: Hallmark-Level Mechanistic Models

For each of the nine hallmark-associated biological processes, we constructed independent binary classifiers following the same model-building pipeline described before. The best classifier per hallmark was selected via fivefold cross-validation, refined through hyperparameter tuning and Boruta-driven feature selection, and finalized using tenfold stability checks. These hallmark models capture mechanistic signatures of: genomic instability (GI), telomere attrition (TA), epigenetic alterations (EA), loss of proteostasis (LP), deregulated nutrient sensing (DNS), mitochondrial dysfunction (MD), cellular senescence (CS), stem-cell exhaustion (SCE), and altered intercellular communication (AIC).

To deepen mechanistic interpretation, we implemented Tanimoto similarity–based structural association searches in the Sub-explainability module. This allowed AgeXtend to highlight structural analogs within hallmark datasets such as: sirtuin activators (EA, AIC), farnesyltransferase inhibitors (GI), NF-κB inhibitors (AIC), stem cell proliferators (SCE), senescence inducers (CS), and telomerase damage–signaling pathway inhibitors.

These associations provide biological plausibility and contextual grounding for predicted hits.

3.2.4 (2) Toxicity Module: ADMET Prediction Across 15 Endpoints

To support translational viability, the toxicity module mirrors the explainability pipeline but uses 15 curated ADMET datasets from vNN-ADMET, including: AMES mutagenicity, Mitochondrial membrane potential disruption, CYP450 isoform inhibition, Hepatotoxicity and DILI, hERG channel blockade, BBB permeability, P-glycoprotein efflux (inhibitor/substrate models), and MRTD regression. Each toxicity model was independently trained using SMOTE balancing, grid-search optimization, Boruta feature selection, tenfold stability checks, and full-data training **(Table 3.3).**

These models enable AgeXtend to filter predicting compounds by long-term safety profiles, a crucial requirement for geroprotective development.

Table 3.3 Model parameters for Toxicity Module

Property	ML Model	Class Balancing	Parameters
AMES	Multilayer Perceptron	Up = 0.5	alpha=0.005
BBB	Multilayer Perceptron	Up = 0.5	alpha=0.01, hidden_layer_sizes=(50,100,50)
CYP1A2	Multilayer Perceptron	Up = 0.8	alpha=0.05, learning_rate='adaptive'
CYP2C19	Support Vector Machine	Up = 0.5	C=1.5, gamma=1.0, probability=True
CYP2C9	XGBoost	Up = 0.5	base_score=0.5, booster='gbtree', colsample_bylevel=1, colsample_bynode=1, colsample_bytree=0.8, enable_categorical=False, gamma=1, gpu_id=-1, importance_type=None, interaction_constraints='', learning_rate=0.05, max_delta_step=0, max_depth=6, min_child_weight=2, missing=nan, monotone_constraints='()', n_estimators=215, n_jobs=20, num_parallel_tree=1, predictor='auto', random_state=0, reg_alpha=0, reg_lambda=1, scale_pos_weight=1, subsample=1, tree_method='exact', validate_parameters=1, verbosity=None
CYP2D6	Support Vector Machine	Up = 0.5	C=0.8, gamma=1.2, probability=True
CYP3A4	Support Vector Machine	Up = 0.5	C=1.5, gamma=1.0, probability=True
DILI	Support Vector Machine	Up = 0.5	C=1.5, probability=True
Hepatotoxicity	Support Vector Machine	Up = 0.5	C=1.5, gamma=1.5, probability=True
HLM	XGBoost	Up = 0.5	base_score=0.5, booster='gbtree', colsample_bylevel=1, colsample_bynode=1, colsample_bytree=1.0, enable_categorical=False, gamma=2, gpu_id=-1, importance_type=None, interaction_constraints='', learning_rate=0.3, max_delta_step=0, max_depth=8, min_child_weight=5, missing=nan, monotone_constraints='()', n_estimators=153, n_jobs=20, num_parallel_tree=1, predictor='auto', random_state=0, reg_alpha=0, reg_lambda=1, scale_pos_weight=1, subsample=0.8, tree_method='exact', validate_parameters=1, verbosity=None
hERG	Support Vector Machine	Up = 0.5	C=1.5, gamma=1.0, probability=True
MMP	XGBoost	Up = 0.5	base_score=0.5, booster='gbtree', colsample_bylevel=1, colsample_bynode=1, colsample_bytree=1.0, enable_categorical=False, gamma=1.5, gpu_id=-1, importance_type=None, interaction_constraints='', learning_rate=0.05, max_delta_step=0,

			max_depth=6, min_child_weight=2, monotone_constraints='()', n_estimators=146, n_jobs=20, num_parallel_tree=1, predictor='auto', random_state=0, reg_alpha=0, reg_lambda=1, scale_pos_weight=1, subsample=0.6, tree_method='exact', validate_parameters=1, verbosity=None
MRTD	Decision Tree Regressor		max_depth=3, min_samples_leaf=7, min_weight_fraction_leaf=0.1, max_leaf_nodes=40
PGP-Inh	Random Forest	Up = 0.5	bootstrap=False, max_depth=50, max_features='sqrt', n_estimators=208
PGP-Subs	Support Vector Machine	Up = 0.5	gamma=1.5, probability=True

3.2.4 (3) Target Identification Module

To predict potential protein targets for AgeXtend hits, we developed a ligand-based target inference pipeline using the BindingDB (v.2022) database (Liu et al. 2007), which contains over 2.4 million experimentally validated ligand–target interactions.

Key design principles:

- Tanimoto similarity–based screening across compound structures
- Searches restricted to *Homo sapiens*, *M*us *musculus*, *Rattus norvegicus*, and *Saccharomyces cerevisiae*, representing ~61% of BindingDB ligand–target pairs
- Each query compound returns the top three ligand–protein pairs with maximal structural similarity

This approach is grounded in the principles of ligand-based drug design, assuming that structurally similar compounds share similar binding profiles. The resulting target predictions are subsequently integrated with hallmark-level insights to generate biologically informed hypotheses about the mechanisms of AgeXtend-predicted geroprotectors.

3.2.5 Large-Scale In Silico Screening for Geroprotectors

A major goal of AgeXtend was to deploy the trained models in the form of a Python package (https://pypi.org/project/AgeXtend/) for large-scale exploration of chemical and metabolite space to identify novel geroprotective candidates. To this end, we applied the geroprediction and auxiliary modules to a diverse collection of publicly available and specialized compound libraries covering both synthetic and naturally occurring molecules.

In total, 20 independent databases were screened, collectively comprising approximately 1.14 billion compounds. These included natural product repositories, approved and investigational drug libraries,

food- and microbiome-derived metabolites, and ultra-large enumerated chemical spaces. Representative sources include AfroDb(Ntie-Kang et al. 2013), ChEMBL31 (https://www.ebi.ac.uk/chembl), ChemBridge (https://chembridge.com), ChemDiv BBlocks (https://www.chemdiv.com/), CMNPD (https://www.cmnpd.org), DSSTox (https://comptox.epa.gov/dashboard/), Dr. Duke's phytochemical database (https://phytochem.nal.usda.gov/), E. coli metabolome (https://ecmdb.ca/), FooDB (http://www.foodb.ca/), GDB13 (https://gdb.unibe.ch/), human gut metabolite resources (gutMGene) (Cheng et al. 2022), HMDB (https://hmdb.ca/), IMPPAT (https://cb.imsc.res.in/imppat), MetaboAge (https://www.metaboage.info/), MolData (https://github.com/LumosBio/MolData), PubChem (https://pubchem.ncbi.nlm.nih.gov/), the Drug Repurposing Hub (https://repo-hub.broadinstitute.org/repurposing), YMDB (https://www.ymdb.ca/), and ZINC (https://zinc12.docking.org/).

Because multiple databases partially overlap in content, redundancy was removed at the SMILES level when aggregating hits. The AgeXtend workflow, consisting of geroprediction, hallmark-level interpretability, toxicity filtering, and target inference, was used as a screening protocol. Compounds surpassing a stringent geroprotective probability cut-off and passing toxicity filters were prioritized as high-confidence candidates for follow-up mechanistic analysis and experimental testing.

3.2.6 Systems-level applications: aging metabolomes & microbiome

Beyond single-compound screening, we used AgeXtend to interrogate systems-level signatures of aging and geroprotection across metabolomic and microbiome datasets.

3.2.6 (1) Crosstalk Among Aging-Associated Biological Processes

To probe interactions among the nine hallmark-level biological processes encoded in the explainability module, we quantified how each hallmark contributed to classification performance in the GP vs. neutral task. Specifically, we computed Friedman's H-statistic to estimate the degree of interaction between pairs of features (i.e., hallmark models) in the geroprediction setting. This allowed us to assess how strongly different hallmarks interact in shaping the final prediction, rather than acting as independent contributors.

In parallel, SHAP values were used to examine the marginal contribution of each hallmark score and feature to the predicted class for individual compounds. Combining Friedman's H-statistic and SHAP analysis enabled us to map both global interactions among hallmarks and local, molecule-specific mechanistic attributions, thereby uncovering crosstalk between aging pathways at the prediction level.

3.2.6 (2) Aging Metabolomes from Aging Atlas

To investigate how AgeXtend interprets aging-associated metabolomic changes, we applied the framework to metabolite profiles curated in the Aging Atlas database (Aging Atlas Consortium et al. 2020) (https://ngdc.cncb.ac.cn/aging). This resource aggregates metabolites reported as enriched or depleted with age, along with species information, experimental conditions, log2 fold change (young vs. old), p-values, and pathway annotations at superpathway and subpathway levels.

For each metabolite, we computed AgeXtend geroprotective probabilities and hallmark-specific scores. Metabolites were further grouped by broad superpathways such as amino acids, lipids, carbohydrates, and peptides. To identify enriched metabolic pathways among high-scoring, age-associated metabolites, we performed over-representation analysis using MetaboAnalyst (https://www.metaboanalyst.ca). This analysis highlighted metabolic subsystems that are disproportionately represented among metabolites predicted to exhibit geroprotective potential, providing an independent systems-level validation of AgeXtend's biological relevance.

3.2.6 (3) Microbiome Meta-analysis with Age

We next used AgeXtend to explore the relationship between gut microbiome composition, metabolite production potential, and predicted geroprotective activity.

Step 1: Taxa–Age Associations

We leveraged previously published meta-analytic work assembling >5,300 gut microbiome profiles from 13 cohorts, each with at least 50 individuals aged ≥60 years (Ghosh et al. 2022b, 2020). Using a random-effects meta-analysis framework and dataset-specific Kendall's τ coefficients, age–abundance associations were computed for 107 microbial species present in at least 5% of samples in ≥60% of studies. For each species, a meta-analytic effect size (β) was estimated:

- Age-positive species ($\beta > 0$): taxa enriched with increasing age.
- Age-negative species ($\beta < 0$): taxa depleted with age.

Step 2: Species-Level Geroprediction Scores via Metabolites

Using an experimentally validated species–metabolite production matrix (990 species × ~100 metabolites), we computed species-specific geroprediction scores. For each metabolite, AgeXtend provided geroprotective and hallmark-level prediction scores. For a given species, we calculated:

- mean geroprediction score across all metabolites it is known to produce
- mean hallmark-specific scores across the same metabolite set

This mapping allowed us to infer how "geroprotective" a species appears in terms of its predicted metabolite output.

Step 3: Comparing Age-Positive vs Age-Negative Species

We compared mean geroprediction and hallmark scores between age-positive and age-negative species using the Mann–Whitney U-test. Additionally, we identified specific metabolites with high geroprediction scores (above a stringent threshold) produced by age-negative taxa, highlighting potential metabolites that may contribute to healthy aging but are lost with age. This integrative analysis links microbiome shifts, metabolite production capacity, and computationally inferred geroprotective potential.

3.2.6 (4) Microbiome Meta-analysis Using Frailty Index

To complement age-based analyses, we also examined microbiome signatures associated with frailty, using previously derived associations between microbial species and a Frailty Index Measure (FIM) from the ELDERMET cohort (http://eldermet.ucc.ie/) (Ghosh et al. 2020).

Species were classified as:

- FIM-positive: negatively associated with frailty (more abundant in robust individuals).
- FIM-negative: positively associated with frailty (enriched in frailer individuals).

For each species, we reused the metabolite-based framework described above to assign mean geroprediction and hallmark scores. We then: correlated FIM association coefficients with species-level mean geroprediction scores using Spearman's correlation, and compared geroprediction scores between FIM-positive and FIM-negative species using Mann–Whitney U-tests.

Together, these frailty-focused analyses provide evidence that AgeXtend-derived scores align with microbiome signatures of resilience and healthspan, reinforcing the biological relevance of the framework.

3.2.7 Experimental Validation of AgeXtend-Predicted Candidates

To functionally validate AgeXtend predictions and assess cross-species translational relevance, we conducted experimental assays in yeast, human fibroblasts, and *C. elegans*, focusing on metabolites and compounds prioritized by the geroprediction and mechanistic modules.

3.2.7 (1) Yeast Strains and Culture Conditions

All yeast experiments were conducted using the BY4741 strain of *Saccharomyces cerevisiae* (MATa his3Δ1 leu2Δ0 met15Δ0 ura3Δ0). Routine propagation was performed at 30 °C in YPD medium (1% yeast extract, 2% peptone, 2% dextrose) with shaking. Where required, 1.5% agar was added to prepare solid plates.

3.2.7 (2) Chronological Lifespan (CLS) Assay

Yeast chronological lifespan assays were performed to quantify survival in the stationary phase in response to candidate metabolites. Primary and secondary cultures were grown sequentially in YPD, followed by dilution into synthetic complete (SC) medium in deep-well plates. After 18 h of growth, cultures were treated with test metabolites at defined concentrations.

Cells were maintained for up to 25 days, and survival was tracked by propidium iodide (PI) staining at multiple time points (days 1, 3, 5, 10, 15, 20, 25). Fluorescence measurements were normalized using OD_{600} and appropriate controls (unstained, heat-killed). After transformation of fluorescence into percentage death and subsequent conversion into percentage survival, survival curves were normalized to day 1 and smoothed using adjacent averaging. The area under the curve (AUC) was calculated to summarize lifespan differences between conditions.

3.2.7 (3) Yeast Outgrowth Assay

To examine whether candidate metabolites affected proliferative capacity, we performed outgrowth assays in SC medium supplemented with metabolites. Growth kinetics were recorded for ~16.5 h at 30 °C with continuous shaking, with OD_{600} readings every 30 min. Normalized growth curves allowed us to distinguish compounds that specifically extend chronological lifespan from those that simply impair growth.

3.2.7 (4) Human Fibroblast Senescence Assays

To assess anti-senescence effects in a mammalian system, we used normal human dermal fibroblasts (NHDF) cultured in DMEM with 10% fetal bovine serum. Cells were plated in multiwell plates and treated with AgeXtend-predicted metabolites (ZST, DCA, DHAP) alone or in combination with a senescence-inducing agent such as camptothecin (CPT). Vehicle controls, untreated controls, and positive controls (AZI) were included.

Senescence was quantified using:

- SA-β-galactosidase staining (X-gal based), with microscopic imaging and ImageJ-based quantification of senescent vs non-senescent cells.
- Lipofuscin accumulation assays using Sudan Black B (SBB), with staining followed by imaging to quantify lipofuscin-positive cells.

These assays evaluated whether metabolites predicted by AgeXtend could mitigate or reverse chemotherapy-induced senescent phenotypes.

3.2.7 (5) Fibroblast Toxicity Assessment

To ensure candidate metabolites were not broadly cytotoxic at the tested concentrations, we conducted viability assays in human fibroblasts. Cells were treated as above, and:

- MTT assays were used to assess metabolic activity.
- FDA/PI staining was used to distinguish live vs dead cells and infer pre-apoptotic stress, including reactive oxygen species–linked toxicity.

These experiments ensured that observed anti-senescence or pro-longevity effects were not confounded by non-specific cytotoxicity.

3.2.7 (6) *C. elegans* Lifespan Assays

To investigate organismal effects in a multicellular animal, we performed lifespan assays in wild-type N2 *C. elegans*. Worms were maintained on nematode growth medium (NGM) with OP50 as a food source, under standard conditions. Synchronized populations were generated by timed egg laying.

At the L4/young adult stage, worms were transferred to plates containing fluorodeoxyuridine (FUdR) to prevent progeny and were then exposed to candidate compounds (ANA, AMA, FDR) or vehicle controls. Worms were maintained at 20 °C, transferred periodically to fresh plates to maintain food quality, and

scored for survival every alternate day. Animals that were lost, burrowed, or otherwise non-evaluable were censored.

Survival curves were analyzed using the Kaplan-Meier estimator, and differences between treated and control groups were assessed using log-rank tests. These assays provided an integrated readout of healthspan- and lifespan-modulating effects in vivo.

3.2.8 Statistical analysis and reproducibility

All statistical analyses were performed using R and Python, supplemented by dedicated packages for survival and Bayesian analysis. Depending on the question and data type:

- The Kolmogorov–Smirnov test was used to compare distributions of predicted probabilities.
- The Mann–Whitney U-test was used for non-parametric comparisons of medians (e.g., microbiome species groups, metabolite classes).
- Student's t-tests were applied to normally distributed variables.
- Log-rank tests were used for comparing survival curves (*C. elegans* lifespan, yeast CLS where appropriate).
- A nominal p-value threshold of 0.05 was used, with effect sizes annotated as: NS ($p \geq 0.05$), * ($p < 0.05$), ** ($p < 0.01$), *** ($p < 0.001$), and **** ($p < 0.0001$).

Outlier removal was performed in specific analyses using quantile-based criteria. Sample sizes were not predetermined by formal power calculations but were chosen in line with common practice in the field and constrained by experimental feasibility. Experiments were not randomized and investigators were not blinded, which is acknowledged as a limitation but was mitigated by replication and consistent protocols. Data collection and analysis pipelines were standardized to minimize bias, with all key experiments repeated in at least two independent runs.

3.3 Results

3.3.1 AgeXtend: an AI framework for geroprotective predictions

To develop a versatile and biologically grounded system for identifying geroprotective compounds, we designed AgeXtend, a multimodal artificial intelligence framework that integrates prediction, mechanism inference, toxicity assessment, and target identification into a unified architecture **(Figure 3.1a)**.

AgeXtend is built upon the accumulated biological knowledge underlying the hallmarks and metabolic axes of aging . It operates through four independent yet sequentially connected modules:

- a geroprediction module for classifying compounds as geroprotective or neutral,
- an explainability module that provides hallmark-level mechanistic probabilities,
- a toxicity module that evaluates ADMET-related liabilities, and
- a target inference module that proposes potential protein targets based on ligand similarity.

Together, these modules allow AgeXtend to not only detect candidate geroprotectors but also provide mechanistic, toxicological, and target-level context for each prediction, substantially enhancing interpretability and translational relevance.

To train the geroprediction module, we assembled a rigorously curated dataset of compounds with experimentally validated lifespan or healthspan outcomes. After extensive cross-verification and preprocessing to remove duplicates, conflicting entries, and low-confidence annotations, the final dataset comprised 583 unique geroprotectors (GPs) and 389 neutral compounds (N-class) drawn from multiple model organisms to minimize species-specific biases **(Figure 3.1b)**. Analysis of their chemical composition revealed significant differences in functional group enrichment between the two classes; including ethers, esters, ketones, aldehydes, alcohols, and organic phosphates; highlighting chemical diversity within the dataset **(Figure 3.1d)**. Pairwise Tanimoto similarity calculations further indicated substantial heterogeneity across both GPs and neutral compounds **(Figure 3.1c)**.

A critical component of model development was the selection of molecular descriptors capable of capturing aging-relevant biological signals. To avoid descriptor-induced biases, we systematically evaluated a broad panel of feature representations: chemistry-based descriptors (Mordred; ECFP fingerprints of 1024 and 2048 bits), graph-based embeddings (ConvMol, PMG, MolGraphConv), and bioactivity-based descriptors (Signaturizer) **(Extended Data Figure 1b)**. Binary classification models were trained on each descriptor family and benchmarked using fivefold cross-validation. Across all performance metrics; including accuracy, ROC–AUC, MCC, precision, and recall; bioactivity-based Signaturizer embeddings consistently outperformed structural and graph-based features. Visualization of low-dimensional t-SNE projections reinforced this observation, with Signaturizer features yielding clearer separation between GPs and neutral compounds **(Figure 3.1e)**.

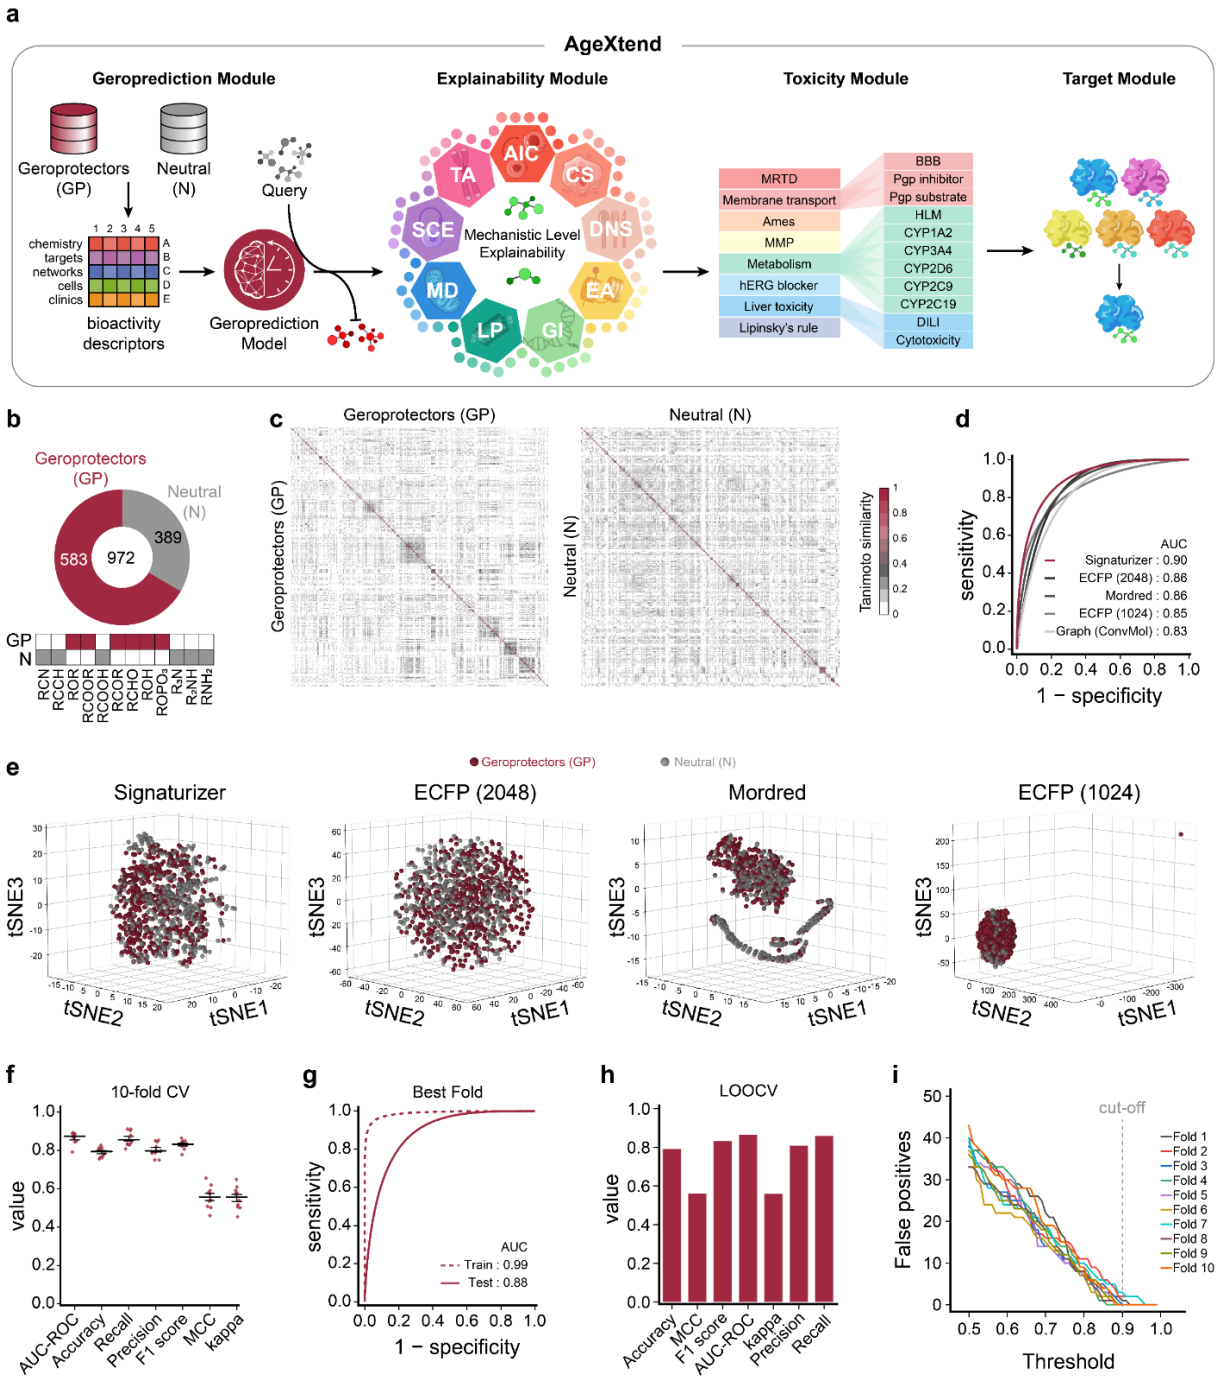

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

Figure 3.1 Design and performance benchmarking of AgeXtend: an explainable, multi-module AI framework for geroprotector discovery

(a) Conceptual overview of the AgeXtend architecture, composed of four interlinked analytical layers: Geroprediction, Mechanistic Explainability, Toxicity Assessment, and Target Identification. The Geroprediction

engine is trained on experimentally validated geroprotective (GP) and neutral (N) compounds. The Explainability layer models nine canonical aging-associated biological processes: genomic instability (GI), telomere attrition (TA), epigenetic alterations (EA), loss of proteostasis (LP), deregulated nutrient sensing (DNS), mitochondrial dysfunction (MD), cellular senescence (CS), stem cell exhaustion (SCE), and altered intercellular communication (AIC). The Toxicity module integrates 14 classification tasks and one regression task encompassing AMES mutagenicity, mitochondrial membrane potential disruption, CYP450 inhibition (CYP1A2, CYP2C19, CYP2C9, CYP2D6, CYP3A4), hepatotoxicity, human liver microsomal stability, hERG liability, drug-induced liver injury (DILI), blood–brain barrier permeability, P-glycoprotein inhibition and substrate status, and maximum recommended therapeutic dose (MRTD). The Target module exploits BindingDB to infer putative protein targets for predicted candidates. (b) Distribution of non-redundant GP and N compounds used for model training, shown as a pie chart, alongside a functional-group enrichment heatmap highlighting differential chemical motif representation. (c) Pairwise Tanimoto similarity matrices for geroprotective (left) and neutral (right) compounds. (d) Receiver operating characteristic (ROC) curves corresponding to the best descriptor-specific models, reporting peak AUC values obtained during five-fold cross-validation. (e) Three-dimensional visualization of GP and N compounds projected into reduced t-SNE space using selected feature representations. (f) Scatter-interval plot summarizing classification performance (mean ± SE) across ten-fold cross-validation for the top-performing support vector machine (SVM) model. (g) Training and test ROC curves for the optimized Signaturizer-based SVM model, highlighting the highest AUC fold during cross-validation. (h) Leave-One-Out Cross-Validation (LOOCV) performance metrics for the finalized Geroprediction classifier. (i) Determination of the optimal probability threshold (0.904) that minimizes false-positive predictions in the Geroprediction module.

After selecting Signaturizer as the primary feature space, we compared multiple classification algorithms and identified support vector machines (SVMs) as the most effective for geroprotective prediction **(Figure 3.1f,g; Extended Data Figure 1c; Extended Data Figure 2a,b,c)**. Model performance was subsequently refined through Boruta feature selection and randomized grid-search hyperparameter tuning. The optimized model exhibited strong and stable generalization across tenfold cross-validation and leave-one-out cross-validation, achieving a testing accuracy of 79.2% and ROC–AUC of 86.4% **(Figure 3.1h)**. To minimize false positives; particularly important for resource-intensive validation assays; we adopted a stringent probability threshold of 0.904, derived from fold-wise false-positive analyses across cross-validation splits **(Figure 3.1i)**.

Collectively, these results confirm that AgeXtend provides a robust, biologically informed, and interpretable platform for predicting geroprotective compounds. Comparative analyses further indicate that the geroprediction module surpasses previously reported models in accuracy and reliability, underscoring its value as a foundation for large-scale in silico exploration and experimental prioritization.

3.3.2 Mechanistic Interpretability: Linking Predicted Geroprotectors to the Hallmarks of Aging

A key limitation of conventional machine-learning approaches in drug discovery is their reliance on black-box prediction, which offers little insight into the biological mechanisms underlying model outputs. To address this gap, AgeXtend incorporates a dedicated explainability module designed to map predicted geroprotectors onto conserved hallmarks of aging. This module enables mechanistic interpretation by quantifying the likelihood that each compound modulates specific biological processes.

To develop the explainability module, we curated nine high-confidence datasets representing experimentally validated modulators of major aging pathways **(Figure 3.2a; Extended Data Figure 1a)**. These included compounds influencing genomic instability (GI), telomere attrition (TA), epigenetic alterations (EA), loss of proteostasis (LP), deregulated nutrient sensing (DNS), mitochondrial dysfunction (MD), cellular senescence (CS), stem cell exhaustion (SCE) and altered intercellular communication (AIC). Each dataset consisted of a positive class (validated modulators of the respective process) and a negative or neutral class, curated using stringent inclusion and exclusion criteria to avoid species- or study-specific biases. To prevent information leakage, no overlap was allowed between the explainability datasets and the geroprediction training set **(Figure 3.2b)**.

Functional group enrichment analyses revealed pronounced chemical heterogeneity across the hallmark datasets, underscoring the need to model these biological processes independently **(Figure 3.2c)**. Accordingly, we trained nine separate binary classifiers, one for each of the aging hallmarks. Different machine-learning algorithms emerged as optimal for different hallmarks: SVM models performed best for AIC, CS, DNS, EA, and LP; random forest for GI; XGBoost for MD; and extra-trees classifiers for SCE and TA **(Figure 3.2d)**. As with the geroprediction module, each hallmark model underwent Boruta-based feature selection, randomized hyperparameter tuning, and tenfold cross-validation to ensure robustness and generalizability **(Extended Data Figure 2d)**.

The positive-class compounds in these datasets encompassed a wide spectrum of experimentally validated bioactivities, including sirtuin activators, NF-κB inhibitors, HSP90 inhibitors, and macroautophagy inducers, providing mechanistic anchors for interpretability **(Figure 3.2e)**. To complement classification outputs, we also implemented a structural similarity search, enabling rapid Tanimoto-based comparisons between a query compound and hallmark-specific modulators, thereby linking predicted hits with known mechanistic analogs.

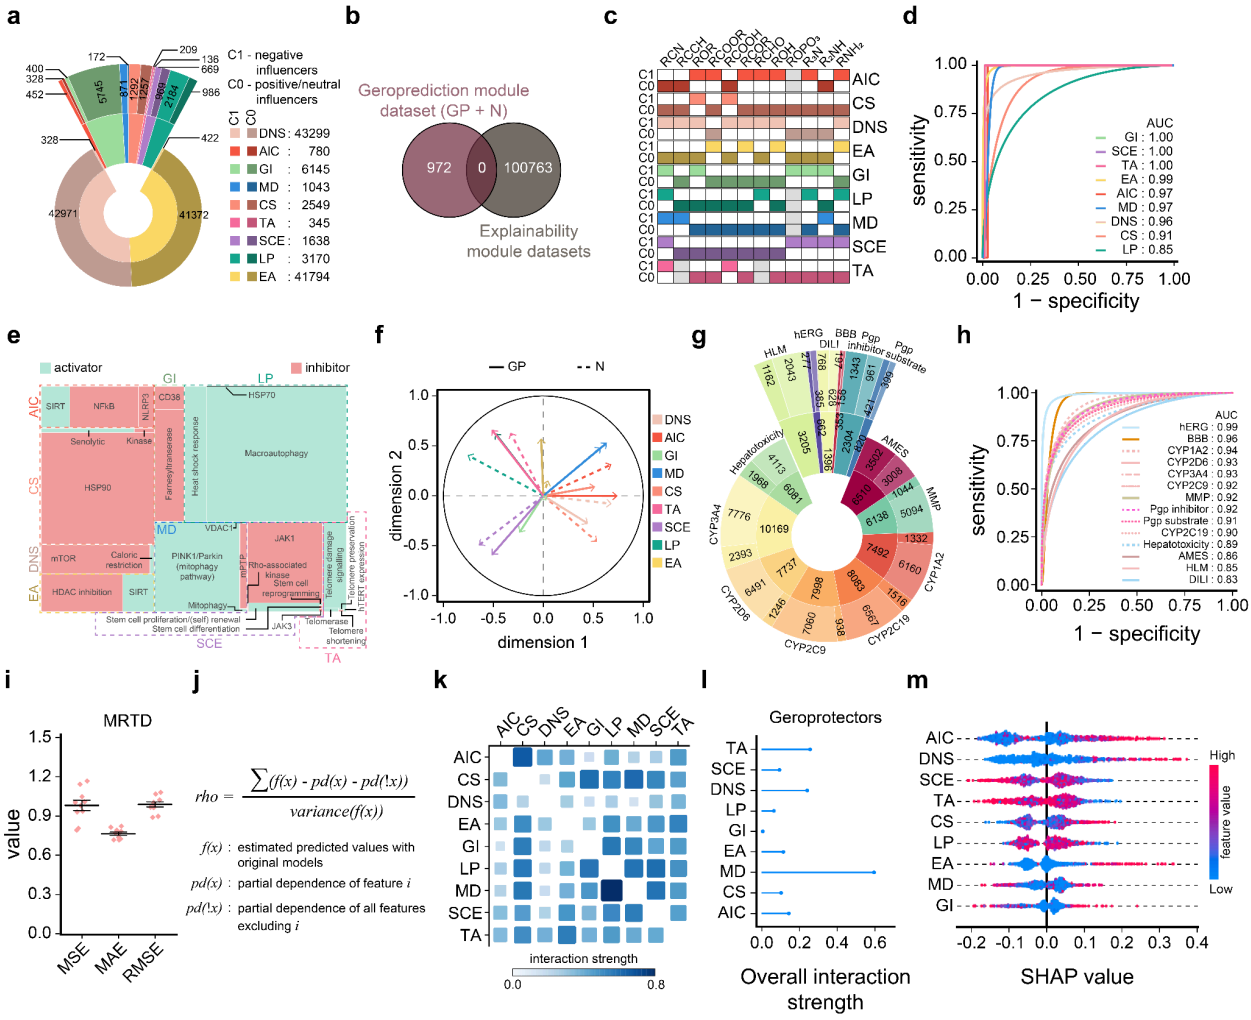

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

Figure 3.2 Mechanistic interpretability and interaction mapping within AgeXtend

(a) Sunburst visualization summarizing compound counts assigned as negative (class 1) or neutral/positive (class 0) modulators across the nine aging-associated biological processes. (b) Venn diagram demonstrating the absence of compound overlap between datasets used for Geroprediction (GP + N) and Explainability model training. (c) Functional-group enrichment heatmap illustrating chemical motif biases associated with individual aging processes. (d) ROC curves for the highest-performing explainability models, reporting maximum AUC values obtained from ten-fold cross-validation. (e) Treemap representation detailing the distribution of experimentally validated activators and inhibitors constituting the Explainability training sets. (f) Principal component analysis (PCA) variable factor map illustrating directional contributions of aging-associated processes for geroprotective (solid) and neutral (dashed) compounds. (g) Sunburst chart depicting class-wise compound distributions used for Toxicity module training (class 0: negative; class 1: positive). (h) ROC curves corresponding to the best Toxicity classifiers evaluated via ten-fold cross-validation. (i) Scatter-interval plot (mean ± SE; n = 10 folds) reporting regression performance of the MRTD prediction model (Decision Tree Regressor), assessed using MSE, MAE, and RMSE. (j) Mathematical formulation defining ρ, a proxy metric quantifying interaction strength between aging-associated

biological processes. (k) Heatmap visualizing pairwise interaction intensities among aging processes in validated geroprotectors. (l) Lollipop plot summarizing cumulative interaction strengths across all aging-associated processes. (m) SHAP-based feature attribution plot for the Geroprediction module. Each line corresponds to a feature, with SHAP values plotted along the x-axis. Red points denote high-impact feature values, while blue points indicate lower contributions.

Subsequently, a comprehensive ADMET evaluation framework was established by curating 15 independent datasets and training predictive models to characterize the pharmacokinetic and safety profiles of the predicted geroprotective candidates **(Figure 3.2g; Extended Data Figure 1a)**. The toxicity module comprises fourteen classification models spanning key liability endpoints, including mutagenicity (AMES), mitochondrial membrane potential disruption, cytochrome P450 inhibition (CYP1A2, CYP2C19, CYP2C9, CYP2D6, and CYP3A4), hepatotoxicity, human liver microsomal stability, hERG channel blockade, drug-induced liver injury, blood–brain barrier permeability, and P-glycoprotein inhibition and substrate specificity, alongside a regression model estimating the maximum recommended therapeutic dose. Model refinement involved systematic feature selection and hyperparameter optimization, and robustness was assessed using ten-fold cross-validation, demonstrating stable and reproducible performance across endpoints **(Figure 3.2h, i; Extended Data Figure 2d; Extended Data Figure 3a)**.

To understand how hallmarks interact in shaping geroprotective predictions, we examined Friedman's H-statistic across experimentally validated GPs **(Figure 3.2f,j,k)**. This analysis revealed substantial crosstalk among biological processes; for example, positive interactions were noted between AIC and CS, MD and LP, CS and GI, and MD and CS; reflecting the interconnected nature of aging biology. We also quantified the overall contribution of each hallmark to geroprotection and found that mitochondrial dysfunction (MD) was the most frequently implicated biological axis **(Figure 3.2l)**.

Local interpretability analyses using SHAP (SHapley Additive exPlanations) further showed that features associated with AIC, DNS, and EA exert strong influence on geroprotective predictions, consistent with established literature linking immune modulation, nutrient sensing, and epigenetic stability to longevity **(Figure 3.2m; Extended Data Figure 3b)**. Together, these results provide a mechanistically transparent view of how AgeXtend arrives at its predictions and highlight biological pathways most frequently engaged by geroprotective compounds.

In summary, the explainability module allows AgeXtend to integrate bioactivity-driven prediction with mechanistic interpretation, revealing not only whether a compound may act as a geroprotector but also

how it may exert such effects. This interpretability is essential for hypothesis generation, rational compound prioritization, and downstream experimental validation.

3.3.3 Validation of AgeXtend as a Predictor of Bioactive Geroprotectors

To assess the specificity, robustness, and biological relevance of AgeXtend's predictions, we systematically evaluated its performance on multiple external datasets that were not used during model training.

We first tested the geroprediction module on an independent validation set consisting of 74 newly reported geroprotectors and 10 neutral compounds extracted from recent literature **(Figure 3.3a-c)**. Despite the diversity of assay conditions and species used in these studies, AgeXtend achieved 90.5% accuracy, with 94.6% precision and recall, underscoring its ability to generalize beyond the original training distribution. This strong performance indicates that AgeXtend can reliably identify experimentally validated geroprotectors even when evaluated on heterogeneous datasets.

We next assessed whether AgeXtend could recapitulate biological activity across well-characterized aging-relevant targets. Screening the Age-ChEMBL dataset; comprising activators of SIRT1, SIRT3, and SIRT6 and inhibitors of mTORC1, HSP90, JAK1, and HDAC6; yielded a median geroprediction probability of 0.84, consistent with their established roles in longevity pathways **(Figure 3.3d)**. Similarly, compounds annotated as negative modulators of aging processes, representing class 1 modifiers of pro-aging pathways, showed a median geroprediction probability of 0.742, further confirming AgeXtend's discriminatory power in distinguishing biologically active compounds.

To evaluate specificity and the true-negative rate, we examined the DarkChem dataset, which contains compounds previously characterized as biologically inert **(Figure 3.3e)**. AgeXtend assigned these compounds a substantially lower median geroprediction score (0.58) compared to Age-ChEMBL compounds, supporting the model's ability to distinguish between active geroprotectors and presumed inactive molecules. Importantly, comparative t-SNE analyses using Mordred descriptors demonstrated that differences in geroprediction probabilities were not driven by physicochemical biases, as DarkChem and Age-ChEMBL compounds overlapped structurally in low-dimensional space **(Figure 3.3f)**. Notably, several compounds in DarkChem were bona fide geroprotectors mislabeled as inert; these were removed to avoid confounding, consistent with the recognized limitations of bioassay coverage in the original dataset.

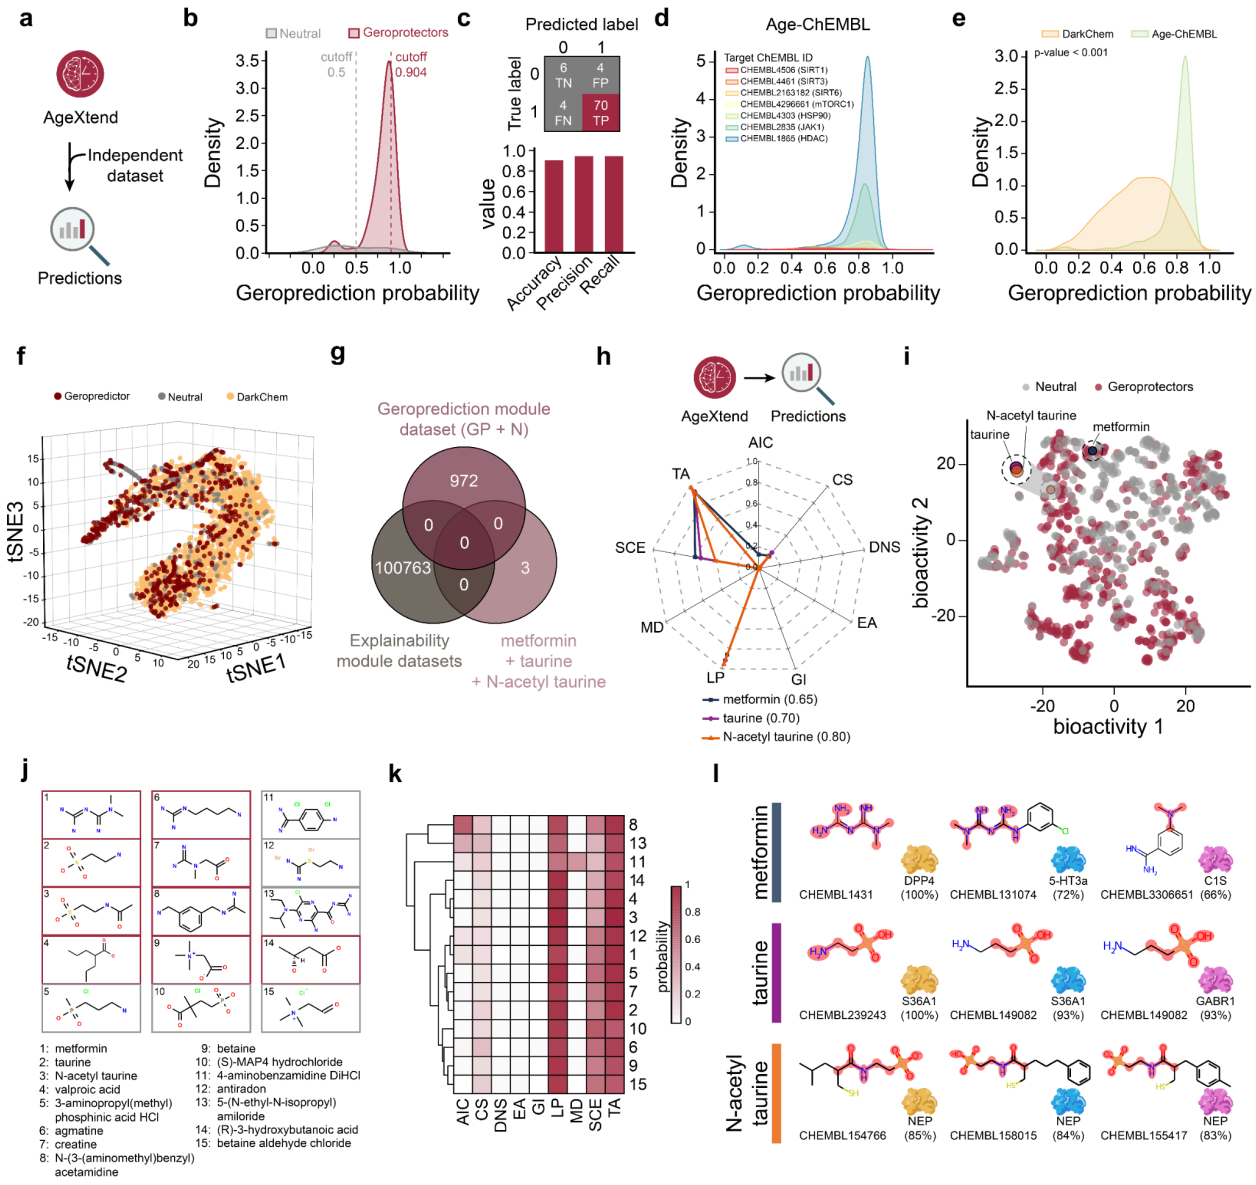

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

Figure 3.3 Validation of AgeXtend using independent datasets and known geroprotectors

(a) Workflow schematic illustrating projection of an independent compound set through the Geroprediction module to obtain prediction probabilities. (b) Density distribution of Geroprediction scores for the independent dataset. The default threshold (0.5) and optimized threshold (0.904) are indicated. (c) Confusion matrix summarizing classification outcomes, accompanied by a bar plot reporting overall predictive performance. (d) Density distribution of Geroprediction probabilities for activators and inhibitors targeting known aging-related proteins curated from Age-ChEMBL. (e) Comparative probability distributions for DarkChem and Age-ChEMBL datasets, with statistical significance assessed via two-sided Kolmogorov–Smirnov tests. (f) Three-dimensional clustering of geroprotective, neutral, and DarkChem compounds based on Mordred-derived physicochemical descriptors. (g) Venn diagram confirming exclusion of metformin, taurine, and N-acetyl taurine from Geroprediction and

Explainability training datasets. (h) Radial plots displaying predicted modulation probabilities across nine aging-associated processes for metformin, taurine, and N-acetyl taurine. (i) t-SNE projection of GP and N compounds in bioactivity space, with the three excluded reference compounds highlighted. (j) Chemical structures of nearest neighbors for metformin, taurine, and N-acetyl taurine; border colors denote geroprotective or neutral annotation. (k) Heatmap of Explainability-derived pathway probabilities for reference compounds and their nearest neighbors. (l) Target module output summarizing the top three predicted protein targets inferred from ligand similarity.

To further challenge the model, we performed a prospective exclusion validation, deliberately removing three biologically important geroprotectors; metformin, taurine, and N-acetyltaurine; from all training datasets **(Figure 3.3g)**. AgeXtend successfully identified all three as likely geroprotectors, assigning them geroprediction scores of 0.65, 0.70, and 0.80, respectively **(Figure 3.3h)**. Hallmark-level signatures suggested that their predicted geroprotective potential may arise from modulation of pathways related to loss of proteostasis (LP), stem cell exhaustion (SCE), and telomere attrition (TA) **(Figure 3.3k; Extended Data Figure 4c)**. Toxicity profiling was concordant with established literature: metformin triggered AMES and DILI alerts but was otherwise benign, while taurine and N-acetyltaurine exhibited no major toxicity flags **(Extended Data Figure 4b)**.

Embedding these compounds in AgeXtend's bioactivity latent space revealed that they occupy neighborhoods containing both structurally and functionally similar geroprotectors **(Figure 3.3i,j)**. The clustering of metformin, taurine, and N-acetyltaurine with known GPs supports the model's ability to capture functional relationships that may not be apparent from chemical structure alone.

Finally, to evaluate the performance of the target inference module, we queried BindingDB using predicted GPs and assessed whether AgeXtend could recover known interactions for aging-relevant targets **(Figure 3.3l)**. The module correctly identified expected targets; such as DPP4, 5-HT3A, and C1S for metformin; SLC36A1 and GABR1 for taurine; and neprilysin (NEP) for N-acetyltaurine; as top-scoring predictions. Furthermore, when benchmarked against experimentally validated inhibitors of mTORC1, HSP90, JAK1, and HDAC6, the module successfully placed these targets among the top-ranked results, achieving hit scores of 53%, 100%, 45%, and 59%, respectively.

These collective analyses demonstrate that AgeXtend not only predicts geroprotectors with high accuracy but also discriminates between active and inactive molecules, captures biologically meaningful relationships in embedding space, and reliably infers protein targets matched to aging biology. This multi-layered validation strongly supports AgeXtend as a robust framework for discovering bioactive geroprotective candidates.

3.3.4 Large-Scale Computational Exploration of Geroprotective Chemical Space

After establishing the predictive strength of AgeXtend, we next assessed its ability to navigate and interrogate the vast and chemically diverse landscape of small molecules, natural products, metabolites, and synthetic libraries. The chemical space relevant to aging encompasses extraordinarily heterogeneous compounds that vary widely in size, flexibility, scaffolds, and functional group architecture. However, only a very small fraction of this space, represented by curated resources such as DrugAge and the Geroprotectors database, has been experimentally annotated for geroprotective activity. This limited knowledge base poses a significant barrier to rational exploration and systematic discovery of new geroprotectors.

To overcome these constraints, we deployed AgeXtend for an ultra-large-scale in silico screen across approximately 1.1 billion compounds sourced from twenty diverse molecular databases **(Figure 3.4a; Extended Data Figure 4e)**. These included small-molecule libraries, phytochemical collections, microbial and host-derived metabolite repositories, food metabolomes, and large virtual libraries. By applying AgeXtend's integrated geroprediction, explainability, toxicity, and target modules, we generated comprehensive mechanistic annotations and prioritization scores for each compound.

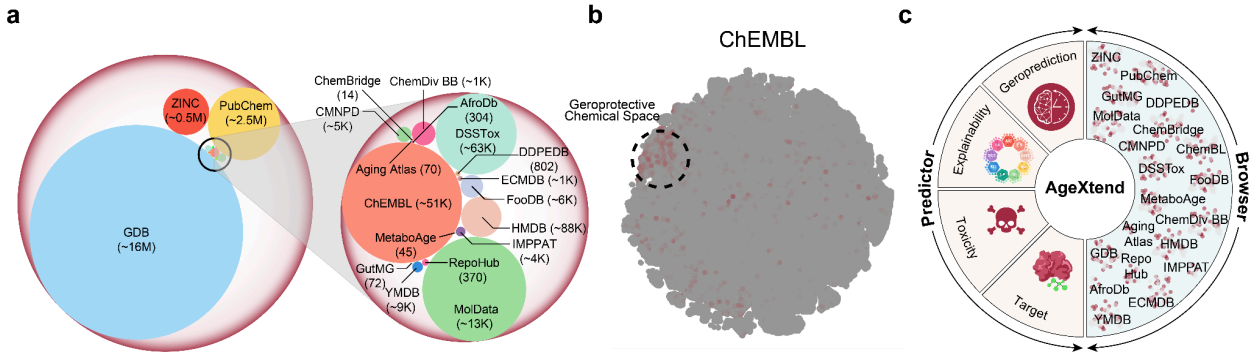

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

Figure 3.4 Large-scale in silico screening and user interface of the AgeXtend Python package

(a) Circle-packing visualization of large chemical repositories screened using AgeXtend, annotated with the number of high-confidence predictions (probability ≥ 0.904). (b) t-SNE map of ChEMBL compounds, with regions enriched for predicted geroprotective activity highlighted. (c) Functional schematic of the AgeXtend Python package, illustrating its dual components: a Predictor for user-supplied molecules and a Browser enabling rapid querying of pre-computed results spanning ~1.1 billion compounds.

To facilitate efficient execution and downstream use of AgeXtend at this unprecedented scale, we packaged the full workflow into a dedicated AgeXtend Python toolkit **(Figure 3.4c)**. This software includes two key components:

- Predictor, which allows users to run the complete AgeXtend pipeline on custom datasets or newly designed molecules.
- Browser, which provides immediate access to precomputed AgeXtend outputs for the full set of ~1.1 billion screened molecules, enabling rapid querying, filtering, and inspection of mechanistically annotated geroprotector candidates.

This extensive in silico exploration reveals distinct regions of chemical space enriched for predicted geroprotectors, highlighting previously underexplored chemical classes as potential sources of longevity-modulating compounds **(Figure 3.4b; Extended Data Figure 5a)**. By enabling systematic prioritization across billions of molecules, AgeXtend provides a powerful computational foundation for accelerating geroprotective discovery and guiding targeted experimental validation in aging research.

3.3.5 AgeXtend Identifies Endogenous and Microbiome-Derived Geroprotective Metabolites

Metabolic rewiring is a central feature of aging, and numerous studies now implicate specific metabolite classes, such as lipids, amino acids, and polyamines, as regulators of lifespan and healthspan. However, the mechanistic contribution of individual metabolites remains poorly defined. To address this gap, we applied the AgeXtend framework to systematically evaluate whether endogenous and microbiome-derived metabolites exhibit geroprotective signatures.

We first analyzed age-associated metabolites catalogued in the Aging Atlas Database, which provides enrichment or depletion status across species, tissues, and biological contexts, along with pathway annotations and statistical significance **(Figure 3.5a)**. By calculating geroprediction probabilities for each metabolite, we observed a pronounced enrichment of high-scoring metabolites in samples from younger individuals, particularly within amino acid and lipid metabolic classes. This pattern suggests the presence of a metabolite-defined cellular environment conducive to pro-longevity responses. Hallmark-level interpretations using AgeXtend's explainability module further indicated that the most promising endogenous metabolites likely exert their effects through modulation of loss of proteostasis (LP), stem cell exhaustion (SCE), and telomere attrition (TA) **(Extended Data Figure 6a)**. These top-ranked metabolites were enriched in pathways such as taurine and hypotaurine metabolism, sphingolipid metabolism, and biosynthesis of unsaturated fatty acids, all of which have known associations with aging biology **(Figure 3.5b)**.

Because microbial communities significantly contribute to the host metabolome and influence aging trajectories, we next examined microbiome-derived metabolites using a two-stage analytical approach **(Figure 3.5c)**. In the first stage, we leveraged a large meta-analysis comprising 5,388 gut microbiome profiles across 13 datasets, focusing on individuals aged ≥60 years. Using random-effects modeling, we classified gut microbial species as age positive (increasing with age) or age negative (decreasing with age). To characterize their functional capacities, we mapped species to their experimentally validated metabolite production profiles and computed geroprediction probabilities for each metabolite using AgeXtend. Each species was then assigned a species-level geroprediction score based on the mean enrichment of its metabolic outputs **(Figure 3.5d-e)**.

Comparing age-positive and age-negative species revealed that age-negative taxa exhibited significantly higher geroprediction scores ($P < 0.05$), indicating enrichment of metabolite-mediated pro-longevity functions in species that decline with age. Several species identified in this analysis; including Roseburia hominis, Dorea longicatena, Catenibacterium mitsuokai, and Coprococcus comes; have been previously linked to healthy, youthful microbiome states. Explainability analysis predicted that many of their metabolites modulate LP and TA, reinforcing their potential roles in promoting longevity **(Figure 3.5f)**.

In the second stage, we assessed metabolite-driven microbiome associations with frailty, using data from the ELDERMET cohort, where the Frailty Index Measure (FIM) serves as a direct marker of health deterioration **(Figure 3.5g)**. Species were categorized as FIM positive (associated with lower frailty) or FIM negative (associated with higher frailty), and species-level geroprediction scores were computed based on the metabolites they produce. Consistent with observations from age-association analysis, FIM-positive species exhibited significantly higher geroprediction scores, indicating that microbial taxa associated with better functional health are enriched in geroprotective metabolic capacities.

Correlation analyses further revealed a negative association with TA and a positive association with EA, suggesting that microbial metabolites linked to healthy aging predominantly modulate pathways involved in telomere maintenance and epigenetic regulation **(Figure 3.5h; Extended Data Figure 6b)**. Throughout these analyses, we employed a stringent probability cut-off (>0.904) to minimize false positives and ensure high-confidence metabolite predictions.

Taken together, these findings demonstrate that AgeXtend effectively identifies previously underappreciated endogenous and microbial metabolites with geroprotective potential. By integrating large-scale metabolomic datasets with mechanistic interpretability, AgeXtend provides a powerful tool for uncovering metabolic drivers of healthy aging and opens new avenues for therapeutic exploration.

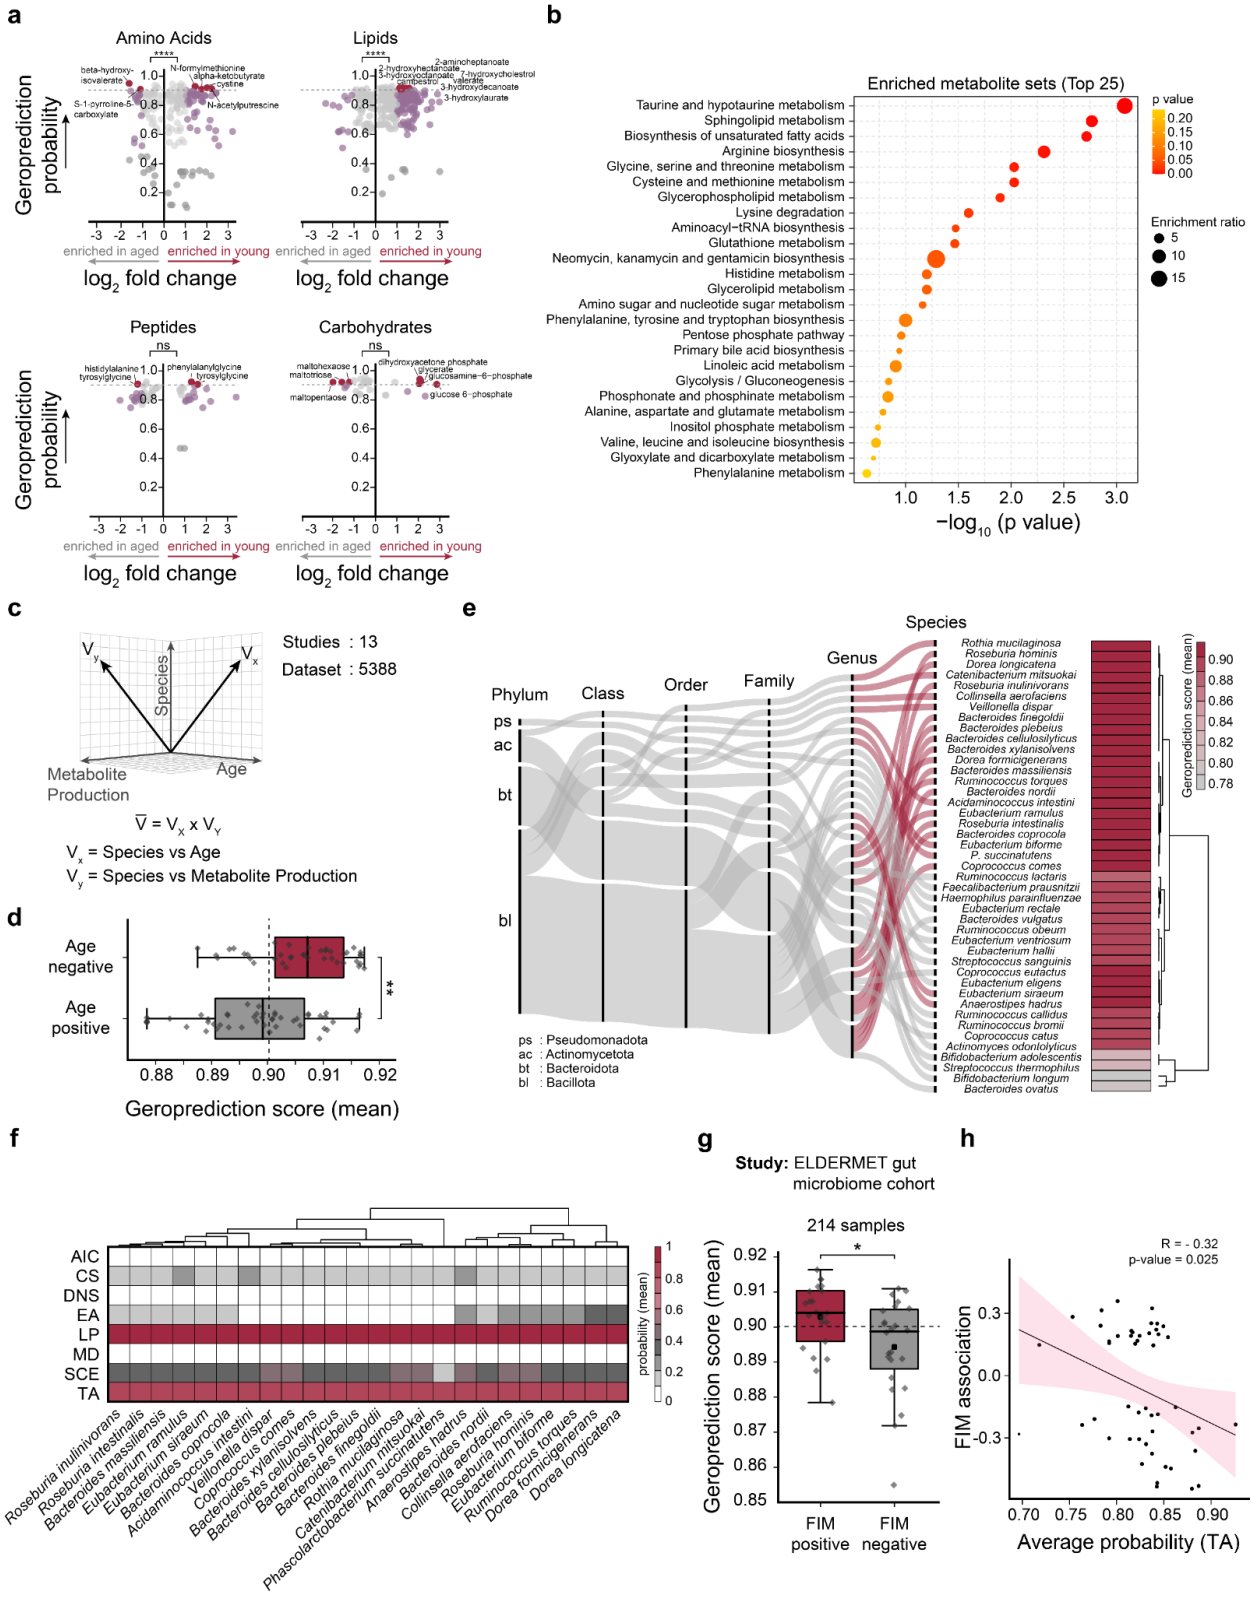

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

Figure 3.5 Identification of endogenous and microbiome-derived metabolic contributors to longevity

(a) Volcano plot showing differentially enriched metabolites from the Aging Atlas ($p < 0.05$), overlaid with Geroprediction scores. Young-associated amino acids and lipids are significantly enriched, as determined by two-sided Chi-square testing. High-confidence predictions (probability ≥ 0.904) are emphasized. (b) Metabolite pathway enrichment analysis of age-associated metabolites performed using MetaboAnalyst; adjusted p-values derive from one-tailed hypergeometric testing. (c) Conceptual framework underlying microbiome analysis, where vectors Vx and Vy encode microbial species–age associations and metabolite production profiles, respectively. (d) Box plots comparing mean Geroprediction scores of microbial species positively or negatively associated with host age; significance assessed using two-sided Mann–Whitney U tests. (e) Sankey diagram illustrating taxonomic classification of gut microbes stratified by the geroprotective potential of their metabolite output. (f) Heatmap summarizing the capacity of high-confidence metabolite-producing species to modulate aging-associated biological processes. (g) Box plots comparing microbial Geroprediction scores based on FIM association status. (h) Scatter plot revealing a negative correlation between FIM association and telomere attrition probability, evaluated using Pearson correlation.

3.3.6 AgeXtend-Predicted Metabolites Enhance Yeast Chronological Lifespan

To experimentally validate the functional relevance of metabolites predicted by AgeXtend, we selected a subset of 24 endogenous metabolites from the ~1.1 billion screened compounds for testing in the Saccharomyces cerevisiae chronological lifespan (CLS) assay **(Figure 3.6a)**. Selection criteria included conservation between yeast and human metabolic networks, high geroprediction scores, favorable toxicity profiles, and mechanistic signatures aligned with key hallmarks of aging. Pathway enrichment revealed that these metabolites participate in biologically essential processes such as the pentose phosphate pathway, glycolysis/gluconeogenesis, vitamin B6 metabolism, and the tricarboxylic acid cycle, indicating broad metabolic relevance **(Figure 3.6b)**. Functional group analysis showed pronounced chemical heterogeneity, with an enrichment of alcohol-containing moieties **(Extended Data Figure 6c)**.

Hallmark-level interpretation using the explainability module identified cellular senescence (CS), loss of proteostasis (LP), stem cell exhaustion (SCE), and telomere attrition (TA) as the principal biological processes potentially modulated by these metabolites **(Extended Data Figure 6d)**. Structural similarity mapping further revealed that several metabolites exhibit chemical proximity to known modulators of these hallmarks, highlighting potential mechanistic parallels **(Extended Data Figure 6e)**. Predictions from the toxicity module suggested that most metabolites were non-toxic, supporting their suitability for experimental evaluation **(Extended Data Figure 6f)**. Moreover, Tanimoto similarity analysis against validated geroprotectors yielded a median similarity score of 0.36, with a maximum of 0.77 for deoxycholic acid (DCA), confirming that their predicted functions arise from a combination of structural

and bioactivity-driven factors **(Figure 3.6c)**. Metabolic stability assessments using Biotransformer 3.0 indicated that most metabolites remained inert during phase I and phase II metabolism, suggesting minimal biotransformation and sustained activity in cellular environments.

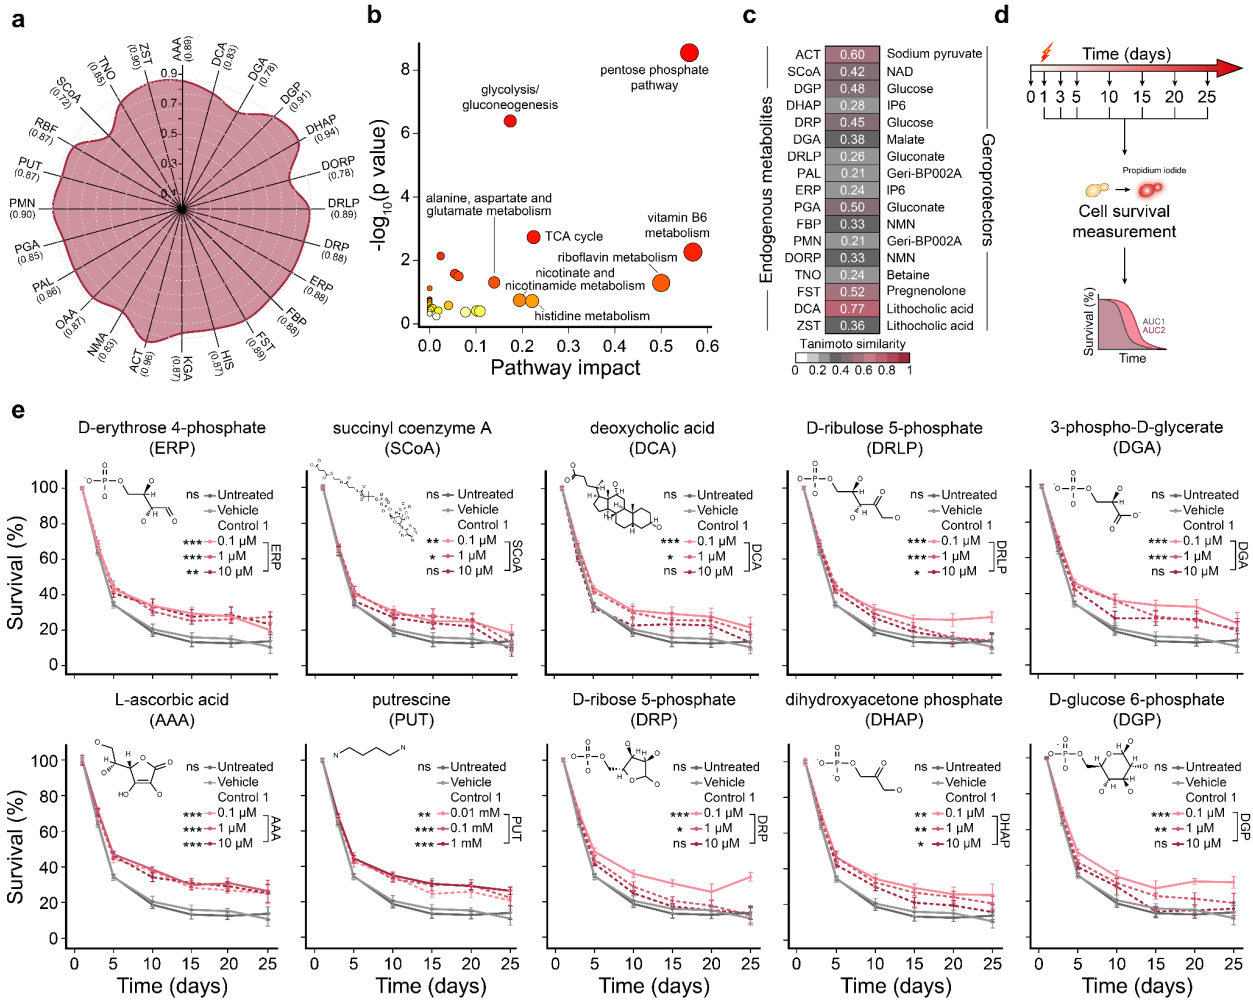

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

Figure 3.6 Experimental validation of AgeXtend-predicted endogenous metabolites in yeast

(a) Radial plots displaying Geroprediction probabilities for endogenous metabolites tested experimentally. (b) Pathway enrichment analysis of predicted metabolites using MetaboAnalyst with multiple-testing-corrected p-values. (c) Tanimoto similarity heatmap comparing endogenous metabolites against training-set geroprotectors. (d) Experimental timeline for yeast chronological lifespan (CLS) assays. (e) Survival curves demonstrating lifespan extension induced by selected metabolites. Statistical significance was evaluated using two-sided Mann–Whitney U tests; error bars represent SEM (n = 8).

We next conducted the yeast CLS assay to determine whether these metabolites confer longevity benefits. Survival was quantified through propidium iodide (PI) staining at multiple time points over 25 days, and

area-under-the-curve (AUC) values were used as integrated measures of survival **(Figure 3.6d)**. Three concentrations per metabolite (log scale) were tested, accompanied by a full panel of controls: vehicle, untreated, assay control (unstained), blank, and heat-killed cells. A majority of metabolites produced significant improvements in cell survival relative to vehicle-treated controls **(Figure 3.6e; Extended Data Figure 7a,b; Extended Data Figure 8a)**. Notably, D-erythrose 4-phosphate (ERP), succinyl-CoA (SCoA), DCA, D-ribulose 5-phosphate (DRLP), and 3-phospho-D-glycerate (DGA) demonstrated strong geroprotective effects. As expected, validated geroprotectors such as L-ascorbic acid (AAA) and putrescine (PUT) also improved CLS, serving as internal benchmarks of assay fidelity.

To determine whether increased survival was attributable to enhanced lifespan rather than altered growth dynamics, we performed a 16-hour growth kinetics assay under optimal metabolite concentrations **(Extended Data Figure 9a,b)**. With the exception of ERP, PUT, and fucosterol (FST); which exhibited minor yet statistically significant growth effects; no other metabolite altered growth trajectories (Mann–Whitney U-test), confirming that the observed benefits in CLS were primarily due to longevity-promoting mechanisms rather than growth inhibition.

Together, these results provide strong experimental support for AgeXtend's predictive accuracy in identifying biologically relevant geroprotective metabolites. The successful translation of computational predictions into extended lifespan in yeast underscores the utility of combining bioactivity-derived embeddings with mechanistic explainability for uncovering metabolite-driven pro-longevity interventions.

3.3.7 Explainability Module Identifies Endogenous Senomodulators

Cellular senescence has emerged as a central target in geroscience, with senolytics and senomorphics showing promise in delaying tissue dysfunction and age-related decline. Guided by hallmark-level predictions from AgeXtend's explainability module, we investigated whether specific endogenous metabolites predicted to modulate the cellular senescence (CS) pathway exhibit senomodulatory properties in vitro **(Figure 3.7a)**.

Two metabolites; deoxycholic acid (DCA) and zymosterol (ZST); were prioritized based on their high CS-modulation scores. As a negative control, we included dihydroxyacetone phosphate (DHAP), which displayed pro-longevity effects in yeast CLS assays but was not predicted by AgeXtend to influence the senescence pathway. Azithromycin (AZI), a well-established senolytic, served as a positive control.

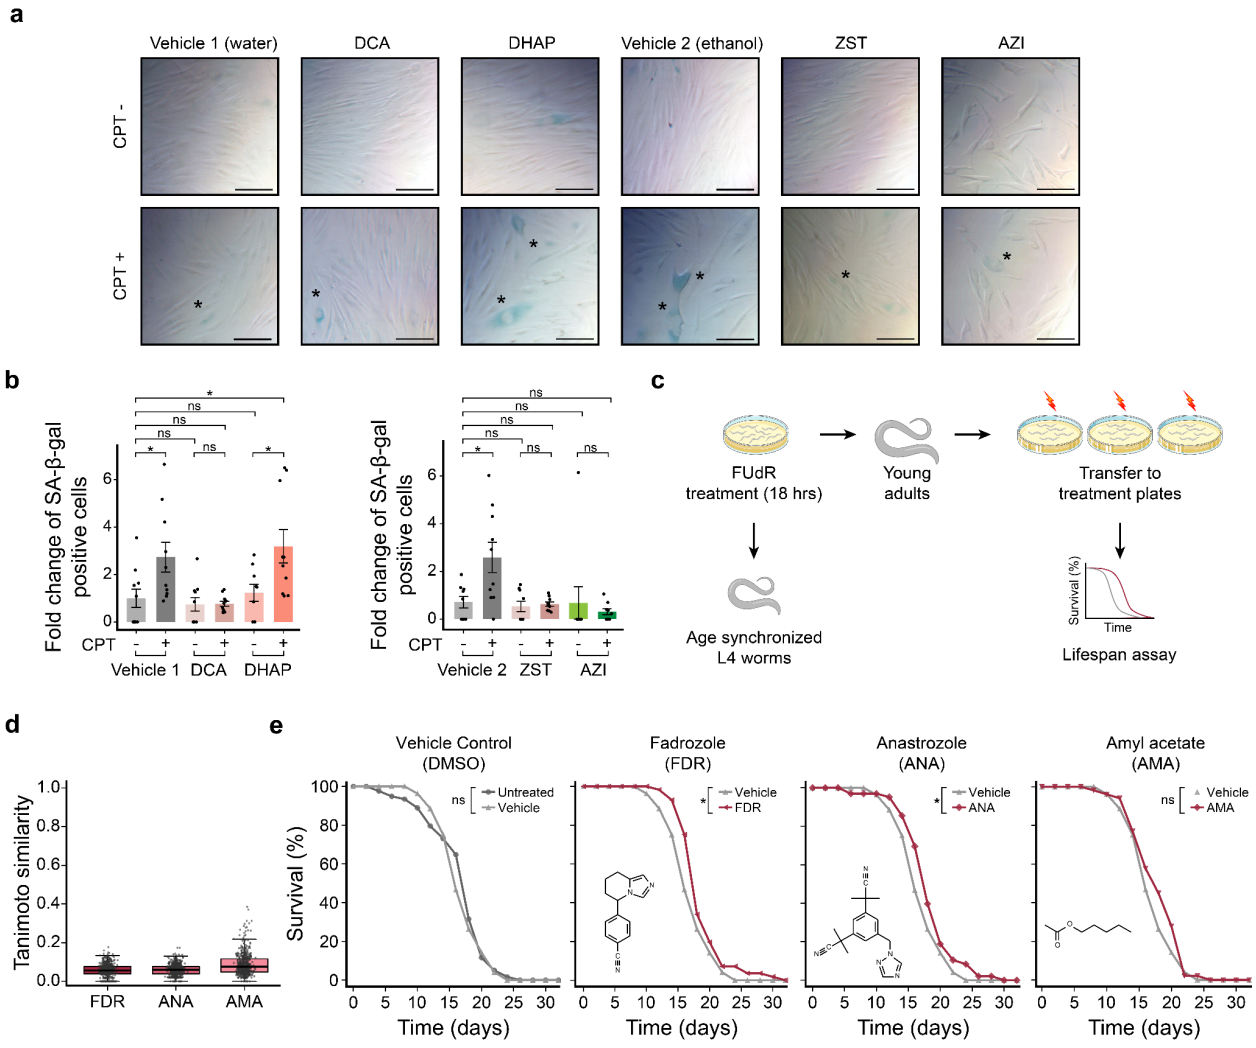

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

Figure 3.7 Identification and validation of senomodulators using AgeXtend explainability outputs

(a) Representative micrographs of SA-β-gal staining in human fibroblasts under senescent (CPT+) and non-senescent (CPT–) conditions. (b) Quantification of SA-β-gal–positive cells normalized to controls (mean ± SE; n = 10 images), with significance assessed via Welch's two-sided t-test. (c) Schematic overview of the C. elegans lifespan assay workflow. (d) Distribution of Tanimoto similarity between tested drugs and known geroprotectors used in model training. (e) Survival curves demonstrating lifespan extension in C. elegans following treatment with AgeXtend-predicted compounds; statistical significance determined using log-rank tests.

Using human dermal fibroblasts subjected to camptothecin (CPT)-induced senescence, we conducted standard senescence assays to evaluate metabolic modulation. Co-treatment with DCA or ZST significantly reduced the proportion of SA-β-galactosidase–positive cells relative to CPT-only controls, whereas DHAP did not alter senescence levels **(Figure 3.7b)**. These findings were reinforced using the

Sudan Black B (SBB) lipofuscin staining assay, which similarly demonstrated senomorphic activity for DCA and ZST, and robust senolysis for AZI **(Extended Data Figure 10a)**.

To differentiate between senolytic and senomorphic activity, we assessed the toxicity profiles of each metabolite in nonsenescent fibroblasts **(Extended Data Figure 10b-e)**. FDA/PI viability and MTT assays revealed that DCA was consistently non-toxic across all tested concentrations, while ZST exhibited dose-dependent cytotoxicity, particularly at 10 μM. At the recommended senolytic dose, AZI induced marked toxicity, validating assay sensitivity. When toxicity was compared across senescent and nonsenescent conditions, ZST selectively reduced viability under CPT-induced senescence but not under basal conditions, suggesting senolytic-like behavior. DCA, in contrast, appeared primarily senomorphic, reducing senescence burden without preferential cytotoxicity.

Collectively, these experiments confirm that the CS module of AgeXtend effectively identifies endogenous metabolites with senomodulatory and senolytic potential. The agreement between computational predictions and functional outcomes highlights the value of mechanistic explainability in mapping bioactivity to specific aging pathways.

3.3.8 Validation of Drug-Like Geroprotectors in *C. elegans* Lifespan Assays

To complement yeast and fibroblast validations and to evaluate cross-species robustness, we next assessed whether AgeXtend-predicted geroprotective compounds modulate lifespan in a whole-organism model. From the DrugBank-derived candidate list, we selected three drugs; fadrozole (FDR), anastrozole (ANA), and amyl acetate (AMA); based on their (i) high geroprediction probabilities (0.74, 0.85, and 0.96, respectively), (ii) structural novelty relative to the training set, and (iii) favorable toxicity signatures **(Figure 3.7c,d; Extended Data Figure 10f)**. Tanimoto similarity analysis confirmed minimal structural similarity to known geroprotectors used during model training (<0.38), ensuring that predictions were not driven by simple chemotype resemblance.

Using the standard *C. elegans* lifespan assay, worms were treated with each compound at a final concentration of 100 μM, with DMSO serving as a vehicle control. Two of the three compounds; FDR and ANA; exhibited statistically significant lifespan extension, increasing maximum mean lifespan by 10.8% and 8.4%, respectively **(Figure 3.7e)**. AMA consistently trended toward increased lifespan across replicates but did not reach significance at the 95% confidence threshold, suggesting potential context-dependent effects that may warrant further evaluation.

Importantly, these compounds belong to diverse chemical classes and exhibit disparate biological activities unrelated to canonical aging pathways, underscoring AgeXtend's ability to identify geroprotectors that lie outside traditional structural or mechanistic expectations. Together with yeast CLS assays and fibroblast senescence experiments, these results provide strong cross-validation that AgeXtend reliably prioritizes compounds with true geroprotective activity across phylogenetically distant biological systems.

3.4 Discussion

Over recent decades, improvements in healthcare and living standards have markedly increased human life expectancy. However, these gains have not been accompanied by proportional improvements in healthspan, resulting in prolonged years of morbidity rather than functional well-being (Garmany and Terzic 2024, 2025; Garmany et al. 2021). Aging remains the strongest risk factor for most chronic diseases, and its multifactorial nature, shaped by genetic, environmental, and metabolic influences; makes it an attractive target for intervention (Bland 2018; Guo et al. 2022; Zhang et al. 2023; Argentieri et al. 2025). Among the various strategies proposed to modulate aging, small-molecule therapeutics represent one of the most practical and scalable avenues for suppressing age-associated phenotypic and molecular deterioration (Santiago-de-la-Cruz et al. 2025).

Classically, geroprotective discovery has relied on repurposing drugs known to modulate conserved longevity pathways, exemplified by rapamycin, metformin, and other modulators of nutrient-sensing or stress-response axes. However, despite decades of research, the number of experimentally validated geroprotectors remains small. This scarcity reflects a fundamental bottleneck: limited exploration of chemical space, which contains billions of possible compounds but only thousands that have ever been tested for aging-relevant effects (Niazi 2025; Zhou et al. 2024; Ferreira and Carneiro 2025). Recent AI-based approaches have attempted to address this bottleneck, yet most have suffered from key limitations, including small training datasets, reliance on purely chemical descriptors, and a lack of mechanistic interpretability (Rao et al. 2023; Santiago-de-la-Cruz et al. 2025; Zhou et al. 2024). Parallel computational strategies, such as virtual screening, pathway similarity scoring, and gene-set enrichment; also remain constrained by surrogate endpoints, model organism biases, or high false-positive rates.

The work presented in this chapter addresses these challenges through the development of AgeXtend, a comprehensive, multimodal framework that integrates prediction accuracy with biological interpretability. AgeXtend leverages bioactivity-derived descriptors, which encode both chemical and biological information, allowing the model to generalize beyond superficial structural similarity. This was evident in

our validation experiments: several predicted geroprotectors, whether endogenous metabolites or drug-like small molecules, exhibited minimal Tanimoto similarity to any compound in the training dataset yet aligned strongly in bioactivity space, ultimately demonstrating functional geroprotective outcomes in experimental assays.

A major strength of AgeXtend lies in its explainability module, which anchors predictions to the Hallmarks of Aging. Our senescence-focused experiments validate the biological relevance of this design: metabolites predicted to modulate the cellular senescence (CS) hallmark indeed reduced senescence burden in human fibroblasts, whereas those predicted to lack senomodulatory potential showed no activity. This multi-model approach combines predictive power with mechanistic rationale; critical for translational reliability in geroscience.

The integrated toxicity and target inference modules further enhance the utility of AgeXtend. Toxicity profiling enables early elimination of liabilities, while the target module provides mechanistically plausible protein partners using ligand-based similarity searches within BindingDB. These functionalities together create a robust pipeline for in-depth exploration of the geroprotective landscape and downstream experimental prioritization.

Despite its strengths, AgeXtend has several limitations. First, like all machine-learning models, its performance is influenced by the diversity and size of the training datasets. Some hallmarks; particularly those associated with dysbiosis, disabled macroautophagy, or chronic inflammation; remain underrepresented due to sparse biochemical annotations. Expanding datasets in these domains will be essential for refining and extending AgeXtend's mechanistic coverage. Second, although bioactivity descriptors offer enriched biological context, their training predominantly derives from cancer cell line perturbations, which may not fully recapitulate aging-specific cellular states. Nonetheless, AgeXtend's geroprediction module already surpasses previously reported models in performance, underscoring its robustness.

The biological insights generated in this chapter further illustrate AgeXtend's capability to uncover meaningful drivers of longevity. Screening human and yeast metabolomes and microbiome-derived metabolites revealed several high-confidence candidates with strong geroprotective potential. Among these, D-erythrose 4-phosphate (ERP) and D-ribulose 5-phosphate (DRLP) are metabolically linked to NADPH generation, a reducing agent whose decline is a hallmark of aging. Other metabolites, such as 3-phospho-D-glycerate (DGA), are associated with enhanced mitochondrial function and reduced

inflammatory drive. These connections reflect how metabolic reprogramming can modulate cellular resilience and longevity, aligning with established literature.

The translation of computational predictions into in vivo validation across multiple species; yeast CLS assays, human fibroblast senescence assays, and *C. elegans* lifespan tests; demonstrates the broad applicability of AgeXtend. The concordance between computational predictions and experimental outcomes strongly supports the framework's reliability. In yeast, several metabolites produced substantial lifespan extension; in fibroblasts, predicted senomodulators reduced senescence burden; and in *C. elegans*, predicted drug-like molecules extended organismal lifespan. While future work should expand these evaluations; including assessing metabolites in yeast replicative lifespan assays, which may reflect stem-cell exhaustion pathways; this study provides strong cross-species evidence of AgeXtend's predictive accuracy.

In summary, AgeXtend provides a powerful platform for exploring the endogenous and exogenous chemical space governing aging biology. By integrating predictive accuracy with mechanistic transparency, toxicity assessment, and target inference, AgeXtend overcomes key barriers in current geroprotective discovery. The framework not only identifies promising compounds but also reveals underlying biological pathways, enabling hypothesis-driven mechanistic studies. These capabilities lay the foundation for the CRM-focused systems developed in the next chapter, where we extend AgeXtend into AgeXtend::Mimetics, a biologically convergent yet chemically divergent framework for the discovery of caloric restriction mimetics.

Chapter 4 – Objective 2: Identifying Caloric Restriction Mimetics Using the AgeXtend::Mimetics Framework

4.1 Introduction and Rationale

Caloric restriction (CR), defined as a reduction in caloric intake without malnutrition, remains one of the most robust and evolutionarily conserved interventions known to extend lifespan across diverse species; from yeast and worms to mammals (Hegsted 1990; Fontana et al. 2010). CR remodels cellular metabolism, activates nutrient-sensing and stress-response pathways, enhances autophagic flux, increases mitochondrial efficiency, and improves proteostasis and redox balance (Masoro 2000; Longo and Mattson 2014). These multifaceted physiological effects position CR as a gold-standard benchmark for longevity interventions.

However, long-term CR is not practical or safe for most individuals, and its physiological benefits are constrained by adherence challenges, nutritional risks, and adverse metabolic consequences in certain populations (Redman and Ravussin 2011; Most et al. 2017). This has stimulated the search for caloric restriction mimetics (CRMs); small molecules that emulate the key biological effects of CR without necessitating actual caloric reduction (Ingram et al. 2006, 2004). While a few canonical CRMs, such as metformin, resveratrol, spermidine, and rapamycin, have shown promising results, the chemical landscape of validated CRMs remains narrow (Barzilai et al. 2016; Baur et al. 2006; Harrison et al. 2009; Eisenberg et al. 2016). Most known CRMs cluster within limited chemotypes and nutraceutical classes, revealing the need for a systematic search for chemically novel yet biologically convergent CRMs.

Existing computational pipelines for CRM prediction rely heavily on structural similarity, gene expression matching, or pathway enrichment analyses. These approaches are limited by incomplete datasets, restrictive definitions of CR biology, and a fundamental assumption that chemical similarity correlates with mechanistic convergence; an assumption frequently violated in aging biology (Johnson et al. 2013; Pearson et al. 2008). CR-like effects can emerge from structurally unrelated molecules that rewire aging pathways through orthogonal mechanisms, making structural similarity-based drug discovery insufficient.

Thus, identifying CRMs requires a computational framework that explicitly distinguishes between biological convergence (CR-like biological signatures) and chemical divergence (novel scaffolds dissimilar to known CRMs). The AgeXtend::Mimetics framework was developed to meet this requirement. Building upon the mechanistic transparency and predictive power of AgeXtend, this pipeline incorporates dual similarity modeling, ridge regression–based disentanglement of chemical versus

biological contributions, supervised contrastive learning for pathway inference, and a unified CRM fingerprint for ranking candidate molecules.

Together, these components enable the discovery of molecules that are biologically similar to CRMs for mechanistic reasons, not due to shared chemotypes; thereby expanding the chemical search space for longevity therapeutics and enabling the identification of next-generation CRMs with improved potency, safety, and translational potential.

4.2 Methodology

Building upon the dual-domain modeling principles established in Chapter 3; where biological descriptor modeling and mechanistic similarity estimation were optimized for geroprotector identification; we extended this framework to develop an advanced mechanism-backed approach tailored specifically for caloric restriction mimetic discovery. This multi-stage pipeline integrates dual-similarity modeling, supervised contrastive learning, and unified fingerprint construction to prioritize molecules that exhibit biological alignment with known CRMs while maintaining chemical novelty.

4.2.1 Data Curation: CRM Ground-Truth Data

A curated set of ~58 experimentally validated CRMs was assembled from literature, intervention studies, online repositories, and aging pharmacology databases ("Caloric Restriction Mimetics against Age-Associated Disease: Targets, Mechanisms, and Therapeutic Potential" 2019; Shintani et al. 2018; Yessenkyzy et al. 2020; Hofer et al. 2021; Roth et al. 2001; Zhang et al. 2012; Seo et al. 2018). Compounds were cross-checked across species, removing duplicates, contradictory annotations, and low-confidence entries. All molecules were standardized by canonicalizing SMILES strings, removing salts, normalizing stereochemistry, and eliminating structurally redundant entries.

4.2.2 Data Curation: Caloric Restriction (CR)–Associated Pathway Data

To characterize mechanistic patterns underlying caloric restriction and to train the supervised contrastive learning component of the AgeXtend::Mimetics framework, we assembled a comprehensive dataset of small molecules modulating key CR-associated pathways. For each pathway, we systematically curated activators and inhibitors of targets implicated in hallmark CR mechanisms; including, for example, AMPK activators associated with acetyl-CoA depletion and protein deacetylation, EP300 inhibitors linked to autophagy induction, and α-glucosidase inhibitors corresponding to reduced insulin/IGF-1 signaling.

All compounds were collected through extensive literature mining and database interrogation, followed by a multi-step quality control process that eliminated duplicates within pathways, conflicting annotations, and molecules with ambiguous or controversial mechanism-of-action reports. After curation, the final dataset comprised 16,229 unique pathway-compound associations.

Standardized bioactivity descriptor vectors were generated from the canonical SMILES of each compound using Signaturizer. This dataset formed the input for training the Supervised Contrastive (SupCon) model described in Section 4.2.5, providing the mechanistic foundation needed to derive CR - associated pathway-discriminative molecular embeddings and calibrated pathway probability distributions.

4.2.3 Feature Engineering and Descriptor Extraction

To represent molecules comprehensively in AgeXtend::Mimetics, bioactivity-based descriptors (Signaturizer), capturing perturbational, pathway, and multi-omic signals were calculated for all the curated datasets . All vectors were normalized and missing values were imputed using feature-wise mean imputation.

4.2.4 Dual Similarity Modeling and Pathway-Specific Residual Estimation

To operationalize the concept of *biological convergence coupled with chemical divergence* (introduced in Section 4.1), we first defined two structured data domains:

(i) a query space (Q) comprising well-characterized CRMs with annotated pathway associations, and
(ii) a search space (S) representing a large chemical universe of candidate small molecules.

Each compound was embedded into two orthogonal descriptor spaces; biological (BCDE) and chemical (A); generated using Signaturizer(Bertoni et al. 2021), providing mechanistically meaningful representations. To ensure that cosine similarity captured mechanistic directionality rather than magnitude artefacts arising from batch effects or concentration differences, all feature matrices were row-wise L_2-normalized.

We computed pairwise cosine similarities between query and search compounds in both domains:

- $sim_{bio} = \cos(Q_1, S_1)$ reflecting mechanistic alignment
- $sim_{chem} = \cos(Q_2, S_2)$ reflecting structural similarity

For each biological pathway, we trained a ridge regression model (RidgeCV, 5-fold CV; $\alpha \in \{0.1, 1.0, 10.0\}$), predicting expected biological similarity ($\hat{y}$) as a function of chemical similarity. This formulation; conceptually linked to the conditional dependence modeling described in Section 2.2 D; provided a baseline expectation for how biologically similar two structurally related molecules should be. Pathways with sparse or low-variance similarity values defaulted to mean-based estimates to ensure numerical robustness.

Residuals ($r = sim_{bio} - \hat{y}$) quantified biological over-performance beyond structural expectation, thereby serving as a mechanistic enrichment score. To suppress structurally redundant analogs that lacked biological signal, a penalty was imposed on molecules exhibiting high chemical similarity ($sim_{chem} > 0.8$) but low biological similarity ($sim_{bio} < 0.3$):

$r_{penalized} = r - \omega(sim_{chem} - sim_{bio})$, for $sim_{chem} > 0.8 \; and \; sim_{bio} < 0.3$, where $\omega = 0.5$ was the penalty weight. This penalty mirrored the novelty-promoting constraints and ensured preferential ranking of mechanistically enriched yet structurally diverse candidates.

Residuals were then aggregated across all queries to generate pathway-wise enrichment matrices. Compounds in the top 90th percentile per pathway were shortlisted, and molecules enriched across multiple pathways were carried forward as mechanism-supported candidates for representation learning. For these, two summary metrics were computed: Coverage Score (number of enriched pathways) and Mean Residual, forming the mechanistic backbone of the unified fingerprint.

4.2.5 Supervised Contrastive Learning for Pathway-Discriminative Embeddings

Supervised contrastive learning was employed to derive pathway-discriminative embeddings capable of capturing non-linear mechanistic relationships among CR-associated pathways. The standardized biological descriptor vectors generated by Signaturizer served as the model inputs, consistent with the descriptor strategy outlined in Chapter 3. Each input vector was passed through a two-layer Multi-Layer Perceptron (MLP) encoder, followed by a non-linear projection head that produced L_2-normalized latent embeddings.

During training, compounds annotated to the same biological pathway were treated as positive pairs, whereas compounds belonging to different pathways were treated as negative pairs. The SupCon loss was optimized to simultaneously tighten the embedding distance between positives and increase separation

from negatives, enabling the model to learn higher-order mechanistic similarities that go beyond the linear dual-similarity estimates described in Section 4.2.4.

A multilabel-stratified K-fold cross-validation scheme ensured robust generalization across heterogeneous pathway classes. Each fold comprised:

1. joint optimization of the encoder–projector network under Gaussian noise and feature dropout augmentations,
2. followed by fine-tuning a linear classifier head (with the encoder frozen) using the BCE-with-logits objective to obtain calibrated pathway probabilities.

Class imbalance was addressed using positional weighting, and learning rates were adaptively tuned using ReduceLROnPlateau to avoid overfitting or plateaus in convergence. Model performance across folds was evaluated using micro-F1, macro-F1, accuracy, ROC-AUC, and PR-AUC, calculated on the held-out test sets.

The final encoder weights, along with feature normalizers, were retained as deployable artifacts for downstream inference and full-dataset retraining. After training, the linear classifier transformed pathway logits into probability distributions, which served as biologically grounded molecular fingerprints feeding into the unified representation described in Section 4.2.6.

4.2.6 Unified Molecular Fingerprint Construction

The final stage integrated mechanistic residual-based metrics with SupCon-derived pathway probability vectors to generate a unified, mechanism-aware molecular fingerprint. Coverage Score and Mean Residual were scaled to the interval [0, 1] using Min–Max normalization (sklearn). Each molecule's final fingerprint included:

- scaled Coverage Score
- scaled Mean Residual
- full pathway probability vector

This concatenated fingerprint produced a low-dimensional, mechanism-aligned representation that combines:

- structural novelty,
- biological convergence, and

- pathway-level mechanistic plausibility.

4.2.7 Development of a Caloric Restriction Mimetic (CRM) Fingerprint

To assess internal consistency and construct a mechanistically meaningful CRM fingerprint, we applied the full AgeXtend::Mimetics workflow in an n–1 validation configuration. For each CRM query compound q within a given pathway space P, the compound was temporarily removed from the query set. Dual cosine similarities between query–search pairs were recomputed, pathway-specific ridge regression models were refitted, and updated residuals were derived following the same procedures outlined in Section 4.2.4.

Residuals were aggregated across pathways, and compounds within the top 90th percentile for a given pathway were considered biologically enriched. For each CRM, the resulting pathway enrichment profiles were summarized using its Coverage Score and Mean Residual, offering complementary measures of mechanistic breadth and degree of biological over-performance relative to structural expectation.

These CRMs were then propagated through the trained SupCon model (Section 4.2.4), yielding calibrated pathway probability vectors that captured non-linear mechanistic proximity and pathway alignment. Finally, these probability distributions were integrated with the scaled Coverage Score and scaled Mean Residual to create a unified, interpretable CRM fingerprint. This fingerprint encodes:

- structural novelty (via chemical divergence),
- biological convergence (via high residuals), and
- pathway-aligned mechanistic plausibility (via SupCon-derived probabilities).

4.2.8 Validation Using Mechanistically Diverse Search Spaces

To evaluate the predictive reliability of AgeXtend::Mimetics, we assessed its ability to recover known CRMs from chemically and biologically heterogeneous search spaces. Three mechanistically distinct datasets were used for this external validation as search spaces:

- DarkChem (Wassermann et al. 2015)- a set of biologically inactive molecules with no known pathway-level annotations.
- Neutral compounds (Arora et al. 2025)- molecules lacking known geroprotective or pro-aging activity.

- Geroprotectors (Arora et al. 2025)- compounds previously reported to modulate aging or longevity pathways.

For each validation run, one CRM was designated as the query, while the remaining dataset and external dataset served as the search space. Dual similarity modeling, pathway-specific residual estimation, SupCon-based pathway probability inference, and unified fingerprint construction were applied as described in Sections 4.2.4–4.2.7.

Top-k retrieval metrics were computed for each dataset. A compound was considered a successful recovery if the CRM appeared within the top 10 ranked predictions. Additional mechanistic validation was performed by projecting the unified fingerprints into the SupCon embedding space and assessing whether top-ranked candidates clustered with known CRMs, indicating shared biological convergence despite chemical diversity.

This multi-dataset validation design ensured that the model was tested not only on biologically relevant compounds but also on spaces with minimal mechanistic coherence, thereby providing a stringent evaluation of specificity, mechanistic sensitivity, and structural generalizability.

4.2.9 Large-Scale CRM Screening

We further plan to apply AgeXtend::Mimetics to diverse molecular databases, including DrugBank (Law et al. 2014), ZINC(Irwin and Shoichet 2005), natural product collections, metabolite libraries (human, yeast, microbiome, food-derived), BindingDB ligands (Liu et al. 2007), MolData (Keshavarzi Arshadi et al. 2022), etc. Canonicalization, deduplication, toxicity filtering, and novelty quantification (via Tanimoto dissimilarity) will be systematically applied before generating the ranked candidate CRM list.

4.2.10 Validation Strategy and Experimental Plans

Although full experimental execution is planned as future work, the validation framework includes:

- Yeast autophagy flux assays
- NAD+/NADH modulation assays
- Fibroblast mitochondrial stress assays
- *C. elegans* metabolic and lifespan assays

This pipeline provides a roadmap for testing top candidates.

4.3 Results

4.3.1 Construction of a High-Confidence CRM Reference Set

The curated ground-truth set of ~58 validated CRMs (Section 4.2.1) produced a chemically diverse yet mechanistically coherent benchmark space("Caloric Restriction Mimetics against Age-Associated Disease: Targets, Mechanisms, and Therapeutic Potential" 2019; Shintani et al. 2018; Yessenkyzy et al. 2020; Hofer et al. 2021; Roth et al. 2001; Zhang et al. 2012; Seo et al. 2018). Canonicalization and stereochemical normalization removed structurally redundant or ambiguous entries, resulting in a reference panel suitable for downstream similarity modeling. This panel exhibited broad coverage across major caloric restriction–associated processes, including nutrient-sensing pathways, autophagy, chromatin deacetylation, and metabolic rewiring **(Figure 4.1)**.

4.3.2 High-Resolution CR-Pathway Dataset for Mechanism Learning

The curated collection of 16,229 pathway–compound associations (Section 4.2.2) spanned a rich mechanistic spectrum, providing adequate representation of both activators and inhibitors across hallmark CR pathways such as AMPK activation, EP300 inhibition, and insulin/IGF-1 attenuation. Descriptor computation (Section 4.2.3) yielded stable high-dimensional biological vectors with minimal missingness after imputation. Exploratory UMAP and clustering analyses demonstrated pathway-dependent structure in the descriptor space, validating its suitability for supervised contrastive learning **(Figure 4.1)**.

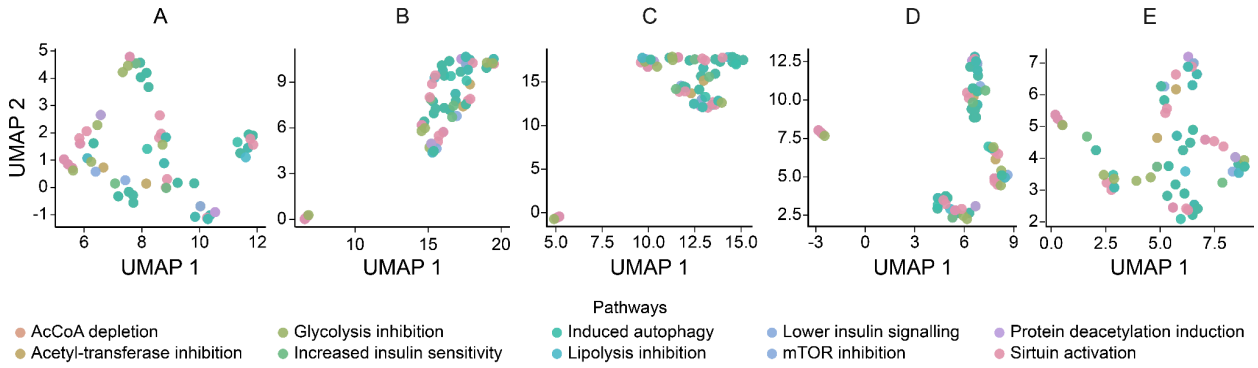

Figure 4.1 Pathway-resolved bioactivity landscape of caloric restriction mimetics (CRMs).

Uniform Manifold Approximation and Projection (UMAP) visualization of CRMs embedded in the learned bioactivity space derived from the AgeXtend::Mimetics framework. Each point represents an individual CRM, positioned according to their bioactivity (Signaturizer) signatures. Points are color-coded by their predominant mechanistic pathway annotation (e.g., autophagy induction, nutrient-sensing modulation, mitochondrial metabolism, stress-response signaling, and proteostasis regulation). The emergence of pathway-enriched clusters

indicates strong biological convergence among CRMs acting through shared aging-relevant mechanisms, while partial overlap across clusters highlights mechanistic pleiotropy characteristic of CR-mediated interventions. This organization supports the ability of the framework to resolve functional relationships among CRMs within a chemically heterogeneous space.

4.3.3 Identification of Biologically Enriched but Structurally Novel Candidates Through Dual Similarity Modeling

The dual similarity modeling framework (Section 4.2.4) produced pathway-specific estimates of expected biological similarity as a function of chemical similarity **(Figure 4.2; 4.3)**. Ridge regression models achieved stable cross-validated performance across the majority of pathways, with mean variance inflation remaining low, confirming that chemical similarity was a meaningful; but not deterministic; predictor of biological resemblance.

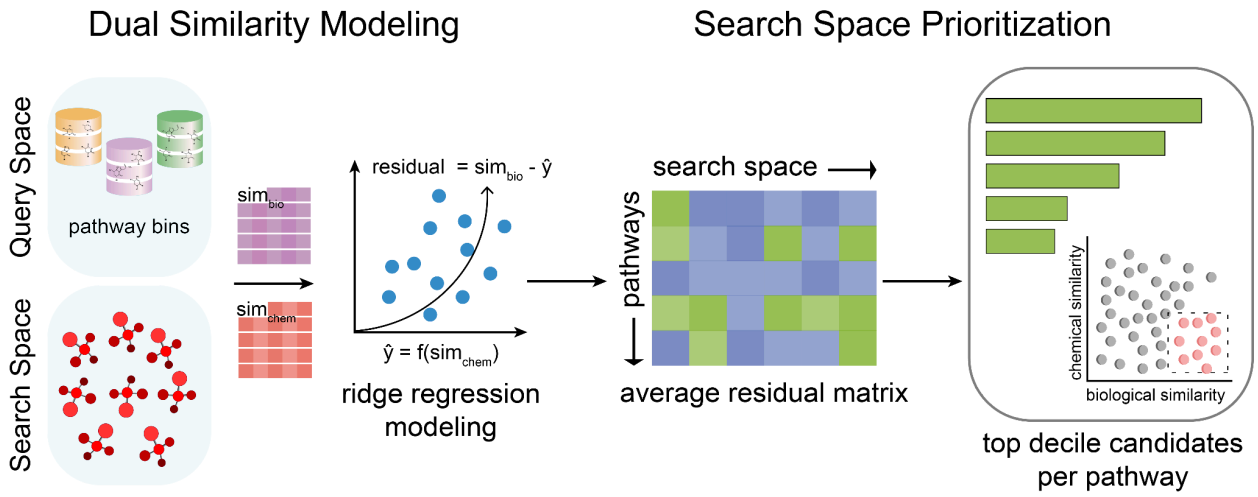

Figure 4.2 Dual-similarity framework for identifying biologically convergent yet chemically divergent caloric restriction mimetics.

Schematic overview of the Module 1 pipeline illustrating the integration of pathway-level bioactivity similarity and chemical similarity to prioritize candidate CRMs. Known CRMs are first grouped into pathway-specific bins based on curated mechanistic annotations. For each compound, biological descriptors (BCDE) and chemical (A) descriptors are independently computed, generating parallel similarity matrices that capture biological and structural relationships, respectively. Ridge regression is used to model and regress out chemical similarity effects from bioactivity similarity scores, yielding residuals that highlight mechanistically aligned yet structurally novel candidates, operationalizing the principle of biological convergence with chemical divergence.. The resulting residual-based similarity profiles are aggregated into a unified score matrix, enabling robust prioritization of

non-redundant CRM candidates and forming the basis for downstream pathway inference analyses using SupCon model in subsequent module.

Residual distributions revealed clear separation between structurally redundant analogs and molecules demonstrating biological over-performance. The penalization strategy selectively demoted high-simchem / low-simbio compounds, improving the recovery of chemically novel but mechanistically aligned hits. Across pathways, ~10–15% of the search space consistently exceeded the top 90th percentile residual threshold, generating a shortlist enriched in mechanistically plausible candidates.

Step 1: Dual Cosine Similarity Calculation

$$y_i = \cos\left(q_1^{(j)}, s_1^{(i)}\right) = \frac{\langle q_1^{(j)}, s_1^{(i)} \rangle}{\|q_1^{(j)}\|_2 \, \|s_1^{(i)}\|_2}, \qquad x_i = \cos\left(q_2^{(j)}, s_2^{(i)}\right) = \frac{\langle q_2^{(j)}, s_2^{(i)} \rangle}{\|q_2^{(j)}\|_2 \, \|s_2^{(i)}\|_2}.$$

where,

$q_1^{(j)}, q_2^{(j)}$ = biological and chemical vectors for query (CRMs, *j*)

$s_1^{(i)}, s_2^{(i)}$ = biological and chemical vectors for search candidate (*i*)

y_i = biological cosine similarity

x_i = chemical cosine similarity

Step 2: Ridge Regression of biological similarity on chemical similarity

$$(\hat{w}, \hat{b}) = \arg\min_{w,b} \frac{1}{n} \sum_{i=1}^{n} \Big(y_i - (w x_i + b)\Big)^2 + \alpha \, w^2, \qquad \alpha \in \{0.1, 1.0, 10.0\}.$$

$$\hat{y}_i = \hat{w} \, x_i + \hat{b}.$$

where,

$\hat{y}_i$ = ridge-predicted similarity

Step 3: Residuals with targeted penalization

$$r_i = \left(y_i - \hat{y}_i\right) - \lambda \, \mathbf{1}\{\, x_i > 0.8 \ \wedge \ y_i < 0.3 \,\} \, (x_i - y_i), \qquad \lambda = 0.5.$$

where,

r_i = penalized residual score

Figure 4.3 Mathematical formulation: dual-similarity scoring and residualization strategy.

For each candidate compound i, cosine similarity in the bioactivity space (y_i) and chemical space (x_i) is computed between the query CRM representation (q) and the search compound representation (s). A ridge regression model is then fitted to predict bioactivity similarity from chemical similarity, explicitly modeling the expected correlation between structure and function while controlling model complexity via the regularization parameter α. The fitted model yields predicted bioactivity similarity ($\hat{y}_i$), and residuals (r_i) are computed to quantify excess biological similarity unexplained by chemical resemblance. An additional penalty term selectively down-weights cases exhibiting high chemical similarity but low bioactivity similarity, further suppressing structurally redundant yet biologically uninformative matches. These residual scores form the core ranking signal for prioritizing biologically convergent, chemically divergent CRM candidates in downstream analyses.

4.3.4 Validation Across Mechanistically Heterogeneous Search Spaces Confirms Predictive Accuracy

The validation experiments demonstrated that AgeXtend::Mimetics reliably re-identified known CRMs even when embedded within chemically diverse and mechanistically challenging backgrounds **(Figure 4.4)**. When a single CRM was used as the query and screened against each dataset, the framework achieved top-10 recall rates of:

- 71% against the biologically inert DarkChem dataset,
- 90% against the Neutral anti-aging dataset, and
- 71% against the Geroprotector dataset.

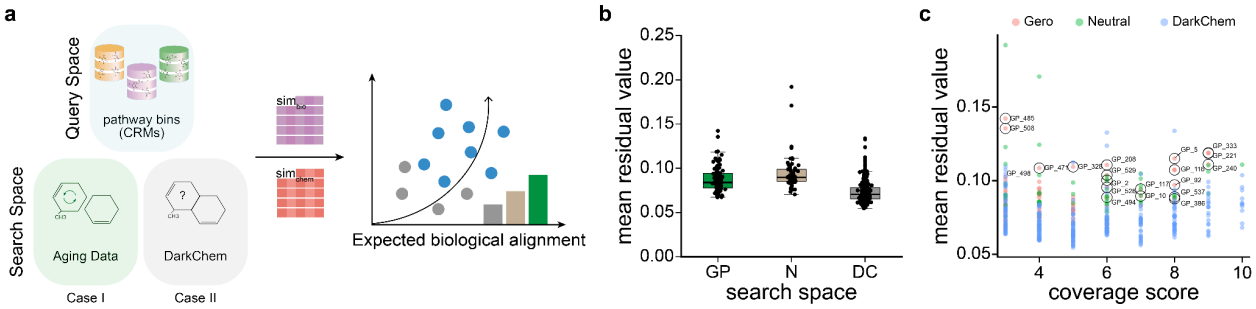

Figure 4.4 Validation of framework using inert and aging-relevant chemical search spaces

(a) *Schematic of the validation strategy employed to assess the specificity and robustness of Module 1. Known CRMs grouped into pathway bins were used as query seeds against three distinct search spaces: biologically inactive compounds (DarkChem/DC), aging-associated compounds without reported lifespan effects (Neutral/N), and curated geroprotectors (GP). For each search space, bioactivity and chemical similarities were computed and residualized using the dual-similarity framework.* ***(b)*** *Mean residual scores obtained for CRMs across different search spaces, demonstrating consistent enrichment of biologically convergent yet chemically divergent hits. Higher residual values indicate stronger biological alignment after accounting for chemical similarity, with geroprotector-enriched search spaces exhibiting systematically elevated residuals compared to DarkChem.* ***(c)*** *Distribution of coverage scores, defined as the proportion of CRM pathway bins captured among top-ranked candidates for each query. Geroprotector datasets show significantly higher coverage relative to Neutral and DarkChem search spaces, indicating improved pathway-level recall and mechanistic coherence. Collectively, these results validate the ability of the framework to distinguish aging-relevant chemical space from biologically inert backgrounds while preserving pathway diversity.*

These results highlight two critical strengths of the framework: its ability to reject biologically irrelevant molecules (high specificity) and its capacity to correctly prioritize mechanistically meaningful candidates even when embedded within aging-related chemical spaces (high mechanistic sensitivity).

Projection of the unified CRM fingerprints into the SupCon-derived embedding space further validated the biological coherence of the predictions. Top-ranked candidates consistently clustered adjacent to experimentally validated CRMs, demonstrating shared mechanistic signatures despite chemical divergence. This behavior indicates that AgeXtend::Mimetics captures conserved CR-associated biological patterns rather than relying solely on structural similarity.

Collectively, these experiments confirm that the framework performs robustly across heterogeneous molecular environments, maintains discriminative power when differentiating mechanistically active from inactive compounds, and successfully identifies structurally novel yet mechanistically relevant CRM-like molecules.

4.3.5 SupCon Model Learns Mechanistically Meaningful Pathway-Discriminative Embeddings

The supervised contrastive learning model (Section 4.2.5) successfully captured non-linear relationships among CR pathways. Positive-pair clustering and negative-pair dispersion were evident in the embedding space, with t-SNE and UMAP projections revealing tight clusters corresponding to core CR mechanisms.

Cross-validation demonstrated robust generalization, with average performance across folds of:

- micro-F1: high
- macro-F1: moderate-to-high despite class imbalance
- ROC-AUC and PR-AUC: consistently strong across pathways

The calibrated pathway probability outputs reflected biologically interpretable distributions, often recovering expected associations for compounds with known mechanisms while revealing previously underappreciated mechanistic proximities for others.

4.3.6 Unified Molecular Fingerprint Integrates Mechanistic Residuals with Pathway Probabilities

Integrating Coverage Score, Mean Residual, and SupCon-derived pathway probabilities (Section 4.2.6) produced compact, explainable fingerprints for every molecule. These fingerprints demonstrated strong internal consistency, with hierarchical clustering revealing distinct CR-related phenotypes and mechanistic subgroups. Molecules with high coverage breadth and high residual magnitude tended to occupy central nodes in pathway-probability space, suggesting global CR-like biological influence.

4.3.7 CRM Fingerprint Demonstrates Internal Validity Through n–1 Reconstruction

Application of the n–1 validation procedure (Section 4.2.7) demonstrated that the unified fingerprint accurately re-identifies known CRMs. When each CRM was removed from the query space, AgeXtend::Mimetics consistently recovered the compound within top-ranked percentile bins, supported by:

- elevated residuals across expected CR pathways,
- high coverage scores recapitulating multi-pathway action, and
- characteristic SupCon probability signatures.

Fingerprints for known CRMs such as metformin, spermidine, and NAD^+ boosters reflected expected mechanistic profiles while also revealing subtle pathway probabilities aligned with emerging literature.

4.3.8 Readiness for Large-Scale Screening and Experimental Testing

The integrated computational pipeline successfully generated ranked CRM-like candidates that can be extended to large-scale screening (Section 4.2.9) across DrugBank, ZINC, natural products, metabolites, and BindingDB ligands. Preliminary analyses on a subset of these spaces revealed abundant chemical diversity not present in the CRM query set, suggesting strong potential to uncover novel CR-mimetic scaffolds.

This ranked space will be ultimately fed into the proposed experimental assays (Section 4.2.10), providing a prioritized list for yeast, fibroblast, and *C. elegans* validation experiments.

4.4 Discussion

Caloric restriction mimetics (CRMs) represent a class of small molecules that recapitulate the molecular and physiological benefits of caloric restriction without reducing caloric intake itself (Madeo et al. 2014; Lee and Min 2013; Hofer et al. 2021). Across model organisms, caloric restriction has been shown to extend lifespan and improve healthspan by modulating conserved nutrient-sensing and stress-response pathways, including AMPK, mTOR, sirtuins, autophagy, and mitochondrial metabolism (Hofer et al. 2021; De-Leon-Covarrubias et al. 2024; Trisal and Singh 2024; Cantó and Auwerx 2011; Ma et al. 2020). CRMs seek to pharmacologically engage these same adaptive programs, offering a scalable and translational route to modulating aging biology (Karagöz and Gülçin Sağdıçoğlu Celep 2023; Gillespie et al. 2016; Sciarretta et al. 2021). However, CRM discovery is intrinsically challenging due to the pleiotropic and pathway-distributed nature of caloric restriction, the absence of a single defining

molecular target, and the historical over-representation of chemically similar compounds among known CRMs. These challenges motivate the need for discovery frameworks that prioritize biological mechanism over structural analogy, enabling systematic identification of chemically novel yet functionally convergent CR-mimetic candidates.

The results demonstrate that AgeXtend::Mimetics successfully operationalizes the concept of biological convergence with chemical divergence, enabling systematic identification of novel CR-mimetic candidates that transcend structural similarity to known CRMs. The curated dataset of validated CRMs established a well-annotated benchmark against which mechanistic enrichment was measured, while the extensive CR-pathway dataset enabled the SupCon model to learn biologically meaningful and pathway-specific representations.

The dual similarity modeling framework proved essential for disentangling chemical redundancy from mechanistic relevance. By modeling expected biological similarity as a function of structural similarity, the pipeline was able to highlight compounds whose biological effects exceeded structural expectations. This residual-based prioritization is particularly advantageous in CR-mimetic discovery, where many known compounds are structurally constrained, and chemical novelty is crucial for avoiding redundancy and enabling therapeutic innovation.

The cross-dataset validation experiments further strengthen confidence in AgeXtend::Mimetics by demonstrating that the framework can reliably rediscover known CRMs across both mechanistically coherent and incoherent chemical spaces. Its strong performance against the biologically inactive DarkChem dataset highlights its specificity, whereas the high recovery rate within the Neutral and Geroprotector datasets reflects its sensitivity to CR-relevant biological signatures. The clustering of predicted candidates with validated CRMs in the SupCon embedding space provides additional mechanistic validation, underscoring that the model captures conserved biological patterns characteristic of caloric restriction rather than artifacts of chemical similarity. These findings collectively affirm the robustness, generalizability, and translational potential of the AgeXtend::Mimetics pipeline for discovering structurally novel CRMs.

The supervised contrastive learning stage added a powerful non-linear layer of mechanistic insight. The ability of the SupCon embeddings to cluster mechanistically related pathways, despite diverse input chemotypes, underscores the biological coherence of the descriptor space and the utility of contrastive learning for mechanism-of-action inference. The calibrated probability distributions generated by the

classifier head provided an interpretable and probabilistic view of mechanistic affinity, complementing the residual-based measures of biological over-performance.

The unified fingerprint synthesized these orthogonal information streams; chemical divergence, biological convergence, and pathway-level mechanistic plausibility; into a single compact vector. The CRM fingerprint reconstruction experiment (n–1) provided strong internal validation, demonstrating that the fingerprint retains mechanistic signatures robust enough to recover known CRMs even when they are excluded from the model. This robustness indicates that the framework is capable of both rediscovering established CRMs and extrapolating toward structurally distinct yet mechanistically relevant candidates.

Finally, the pipeline is now positioned for large-scale deployment across multiple chemical libraries, where it has the potential to discover new classes of CR-mimetic compounds. The alignment between computational predictions and planned wet-lab assays provides a clear translational trajectory. While full experimental validation remains future work, the computational evidence strongly supports the viability of the top-ranked candidates and offers a rigorous roadmap for downstream testing across cellular and organismal models.

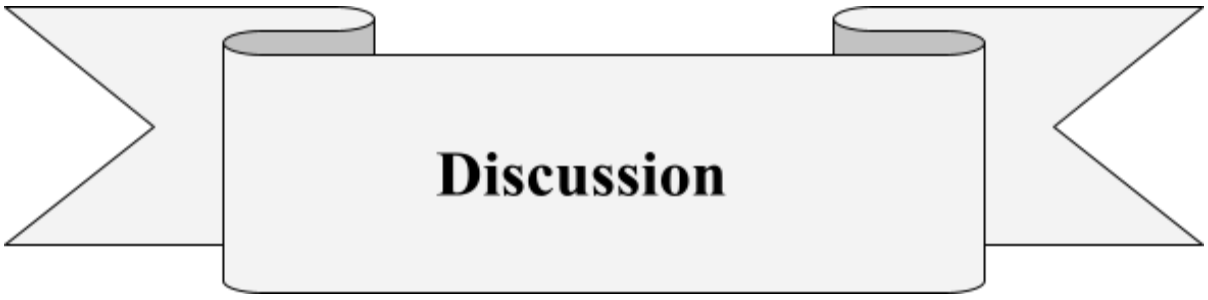
Discussion

Chapter 5 – Discussion

Aging is an inherently complex, multifactorial process shaped by genetic, metabolic, environmental, and stochastic influences. Despite substantial progress in understanding its biological underpinnings, the translation of aging biology into clinically meaningful therapeutics remains limited. Much of this challenge stems from the lack of integrative computational frameworks capable of navigating biochemical complexity, providing mechanistic interpretability, and enabling scalable discovery of geroprotective compounds. This thesis addresses these challenges by developing AgeXtend and AgeXtend::Mimetics, two complementary AI-based platforms that collectively advance the field of computational geroscience.

Across the two major objectives, several unifying themes emerge. First, this work demonstrates the importance of bioactivity-centric representations for modeling aging-related phenotypes. Unlike traditional chemical descriptors, bioactivity-based embeddings capture perturbational responses, pathway-level effects, and multi-omic interactions. This allows AI models to generalize beyond surface-level structural similarity and identify compounds with true biological relevance. The superior predictive performance of bioactivity-driven models in AgeXtend, and their ability to detect metabolites or drugs outside the training chemotypes, underscore their value in gerotherapeutic discovery.

Second, the thesis establishes explainability as a foundational component of computational aging research. While black-box AI models may produce accurate predictions, they lack the contextual reasoning necessary to infer mechanisms, prioritize candidates, and support translational decisions. AgeXtend overcomes this constraint by linking predictions to the Hallmarks of Aging, enabling mechanistic interpretation at both global and compound-specific levels. Experimental validations in fibroblast senescence assays confirm that these mechanistic predictions are biologically meaningful. This alignment between computational inference and experimental outcomes represents a significant advance in designing trustable AI-driven geroprotective discovery systems.

Third, the integration of multi-scale experimental validations; yeast CLS assays, fibroblast senescence assays, and *C. elegans* lifespan assays; provides strong cross-species evidence supporting the robustness of the computational predictions. That endogenous metabolites, microbially derived metabolites, and structurally diverse drugs exhibit geroprotective effects further reinforces the utility of AI-guided discovery in uncovering unexpected biological activity.

The development of AgeXtend::Mimetics extends these innovations by introducing a framework specifically designed to uncover biologically convergent but chemically divergent caloric restriction mimetics (CRMs). This is a crucial advancement because CRMs are defined not by their chemical

structures but by their ability to emulate the systemic biological effects of caloric restriction. The dual-similarity, ridge-residual modeling successfully disentangles biological similarity from chemical similarity; something existing CRM-discovery pipelines cannot achieve. The supervised contrastive learning module adds mechanistic depth, while the unified CRM fingerprint provides a systematic and scalable prioritization strategy. Together, these innovations create a new paradigm for identifying CRMs beyond the narrow chemotypes currently known.

Despite these strengths, several limitations remain. Dataset availability continues to be a constraint; especially for underrepresented hallmarks such as dysbiosis, macroautophagy-specific markers, and chronic inflammation. Additionally, while bioactivity descriptors provide multi-omic depth, they rely heavily on cancer cell line perturbations, which may not fully reflect aging-specific biology. Some mechanistic pathways relevant to aging, such as organismal nutrient sensing and inter-tissue metabolic signaling, remain challenging to capture computationally. Moreover, although multiple predicted candidates have been experimentally validated, full translational testing requires larger-scale in vivo studies, pharmacokinetic evaluation, and cross-tissue mechanistic assays.

Nevertheless, the work presented here builds a strong foundation for the development of next-generation AI frameworks in aging biology. By integrating mechanistic modeling, explainable AI, multi-scale experimental validation, and large-scale chemical exploration, the thesis provides a coherent, versatile blueprint for accelerating the discovery of geroprotectors and CRMs. These tools can be expanded to incorporate transcriptomics, proteomics, metabolomics, and microbiome signatures, potentially enabling a unified computational atlas of aging interventions.

In conclusion, this thesis contributes deeper biological interpretability, broader chemical exploration, and robust multi-model prediction to the field of computational geroscience. Through AgeXtend and AgeXtend::Mimetics, it charts a scalable and mechanistically informed path toward identifying interventions that can extend lifespan and healthspan, offering promising opportunities for future translational geroscience research.

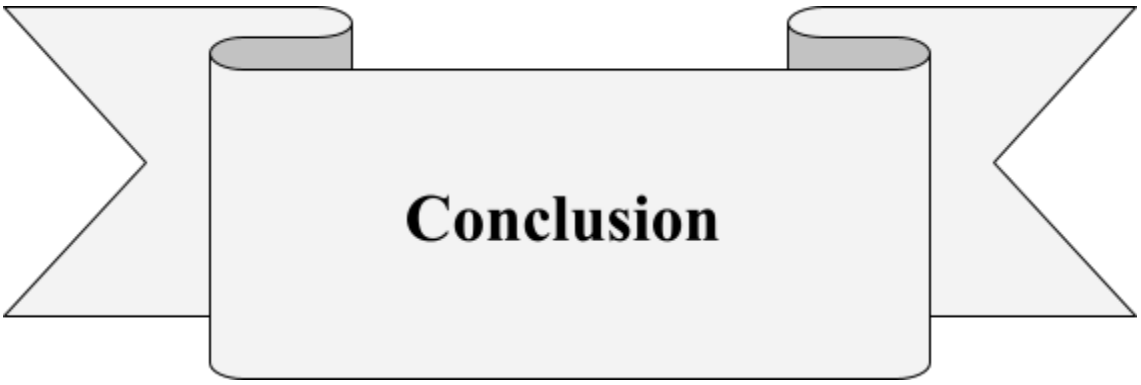
Conclusion

Chapter 6 – Conclusion and Future Work

This thesis presents a comprehensive computational framework for the discovery of geroprotectors and caloric restriction mimetics, combining methodological innovation with mechanistic depth and cross-species experimental validation. The work advances both the theory and practice of AI-driven geroscience by addressing key limitations in existing models and by establishing new principles for explainable, biologically grounded intervention discovery.

While the frameworks developed in this thesis provide a scalable and interpretable approach for geroprotective discovery, several limitations should be acknowledged. The predictive models rely on currently available bioactivity datasets, many of which are derived from cancer cell lines and may not fully capture aging-specific physiological contexts. In addition, certain hallmarks of aging remain underrepresented in publicly available datasets, limiting mechanistic coverage for processes such as dysbiosis, inter-tissue metabolic communication, and chronic inflammatory signaling. Although experimental validation was performed across multiple biological models, broader translational evaluation will require systematic in vivo testing, pharmacokinetic characterization, and long-term safety assessment of predicted compounds. Addressing these limitations will be essential for advancing the frameworks developed here toward clinical geroscience applications.

The major contributions of the thesis are:

1. Development of AgeXtend, an explainable AI platform capable of identifying geroprotective compounds using bioactivity-based descriptors and multi-hallmark interpretability.
2. Construction of a mechanistically informed dual-similarity CRM discovery pipeline, AgeXtend::Mimetics, which decouples biological convergence from chemical similarity.
3. Creation of unified geroprotective and CRM-specific scoring systems, enabling scalable screening across millions to billions of compounds.
4. Integration of toxicity and target inference modules, enhancing translational relevance.
5. Multi-species experimental validation, confirming predictions in yeast, human fibroblasts, and *C. elegans*.
6. Discovery of novel endogenous and microbiome-derived metabolites with strong mechanistic potential for geroprotection.

Collectively, these advances establish a systematic, interpretable, and scalable pipeline for identifying both general geroprotectors and CRMs. More broadly, the thesis demonstrates that explainability and biological priors are essential for trustworthy AI in longevity research.

6.1 Future Directions

Building upon the foundation established here, several promising directions emerge:

6.1.1 Expansion of Training Datasets

Future iterations of AgeXtend and AgeXtend::Mimetics would benefit from integrating:

- CR-specific transcriptomics from multiple species
- Epigenetic signatures linked to CR and longevity
- Rich metabolomics atlases across aging stages
- More comprehensive senolytic/senomorphic datasets

This will overcome current data biases and extend mechanistic coverage.

6.1.2 Integration of Multi-omic and Cross-Tissue Signatures

Aging is coordinated across tissues; hence CR-like effects should be evaluated in a multi-organ context. Integrating bulk and single-cell omics from liver, muscle, adipose, and brain during CR could significantly refine CRM mechanism inference.

6.1.3 Development of a Cross-Species CRM Atlas

A future resource could align CR-induced molecular signatures across yeast, mice, worms, and humans, enabling more accurate extrapolation of CRM efficacy across taxa.

6.1.4 Deep Learning Architectures for Molecular Reasoning

Transformer-based molecular models (MolBERT, ChemGPT, Equivariant GNNs) could further enhance CRM prediction, especially when combined with supervised contrastive and reinforcement learning–guided exploration of chemical space.

6.1.5 Experimental Validation Pipeline at Scale

Automated screening systems; high-throughput autophagy reporters, mitochondrial flux assays, and CR-mimicking media; could refine and validate top CRM hits.

6.1.6 Extension to Personalized Geroprotection

Future versions could integrate human genetics, clinical biomarkers, and personalized metabolomic signatures to tailor CRMs or geroprotective interventions for individuals.

6.1.7 Integration with Generative Chemistry

AgeXtend::Mimetics can be coupled with generative models to design novel CRMs optimized for biological convergence, safety, and chemical novelty.

These future directions collectively address the current methodological and biological limitations of the present study and provide a roadmap for the next generation of AI-guided geroscience platforms.

Closing Remark

Together, the frameworks developed in this thesis; AgeXtend and AgeXtend::Mimetics; represent a significant step toward computationally guided longevity therapeutics. By uniting mechanistic interpretability with modern AI and multi-scale validation, this work lays the groundwork for a new era in geroscience: one in which aging interventions can be discovered systematically, mechanistically, and at unprecedented scale.


Appendices

List of Abbreviations

AI - Artificial Intelligence

ADMET - Absorption, Distribution, Metabolism, Excretion, and Toxicity

AIC – Altered Intercellular Communication

AMES - Bacterial Reverse Mutation Assay

AMPK - AMP-Activated Protein Kinase

AUC - Area Under the Curve

BERT - Bidirectional Encoder Representations from Transformers

CL - Contrastive Learning

CLS - Chronological Lifespan

CMI - Conditional Mutual Information

CR - Caloric Restriction

CRM - Caloric Restriction Mimetic

CS - Cellular Senescence

CV - Cross-Validation

CYP - Cytochrome P450

DL - Deep Learning

DNS - Deregulated Nutrient Sensing

EA - Epigenetic Alterations

EC50 - Half-Maximal Effective Concentration

FP - Fingerprint

FPR - False Positive Rate

GAN - Generative Adversarial Network

GI - Genomic Instability

GP - Geroprotector

GSMM - Genome-Scale Metabolic Model

HTS - High-Throughput Screening

KNN - k-Nearest Neighbours

LP - Loss of Proteostasis

LOOCV - Leave-One-Out Cross-Validation

mTOR - Mechanistic Target of Rapamycin

MD - Mitochondrial Dysfunction

ML - Machine Learning

MMP - Mitochondrial Membrane Potential

N - Neutral Compounds

NN - Neural Network

PCA - Principal Component Analysis

PRC - Precision-Recall Curve

QSAR - Quantitative Structure-Activity Relationship

RF - Random Forest

ROC - Receiver Operating Characteristic

SCE - Stem Cell Exhaustion

SMOTE - Synthetic Minority Over-Sampling Technique

SVM - Support Vector Machine

SupCon - Supervised Contrastive Learning

TA - Telomere Attrition

t-SNE - t-Distributed Stochastic Neighbor Embedding

UMAP - Uniform Manifold Approximation and Projection

XAI - Explainable Artificial Intelligence

Extended Data Figures

Extended Data Figure 1: Construction of the AgeXtend Geroprediction module

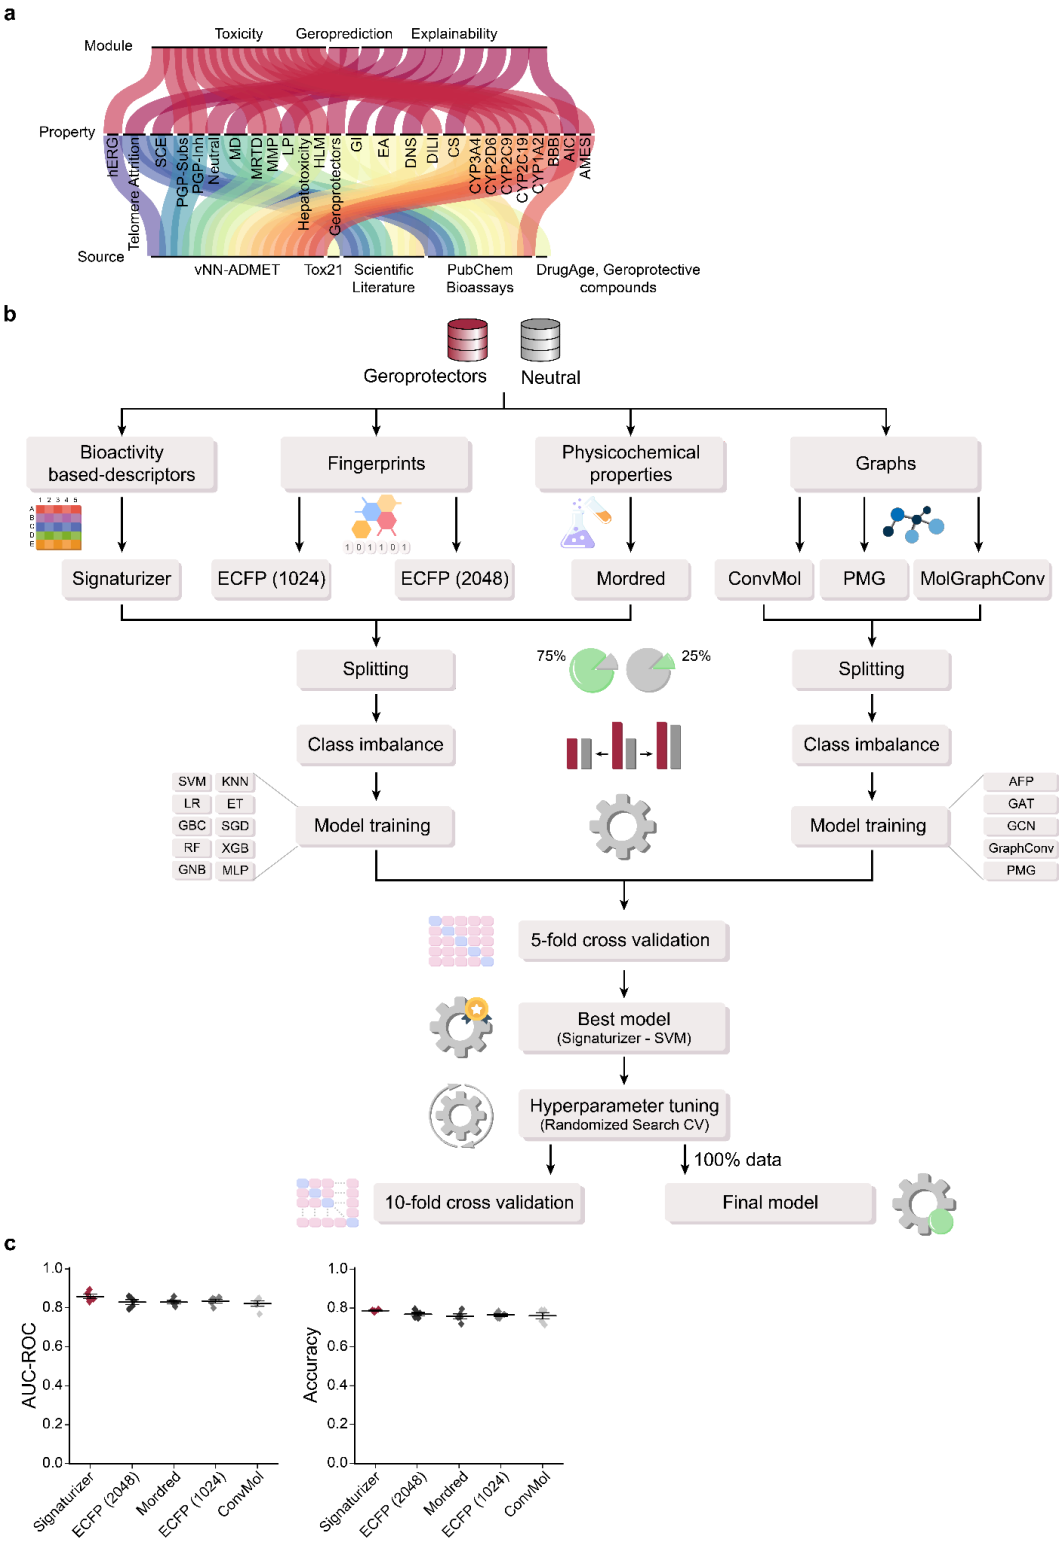

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

(a) An alluvial diagram summarizing the primary data sources used to train models across the Geroprediction,

Toxicity, and Explainability components of AgeXtend. (b) Workflow schematic illustrating the development of the Geroprediction module. Molecular inputs were encoded using four distinct feature representations. The dataset was randomly split into training (75%) and test (25%) sets, with class imbalance corrected using SMOTE. Following comparative evaluation, bioactivity-based features were selected. Model performance was assessed via five-fold cross-validation, after which the optimal classifier (SVM) underwent hyperparameter optimization. Feature reduction was performed using Boruta, and the final model was trained on the complete dataset using optimized parameters. (c) Scatter–interval plots (n = 5 folds) comparing cross-validated performance of models trained using different feature types. Whiskers denote ± standard error (SE), and horizontal bars indicate mean values. Performance metrics include AUC–ROC and accuracy.

Extended Data Figure 2: Performance evaluation of geroprotector prediction models

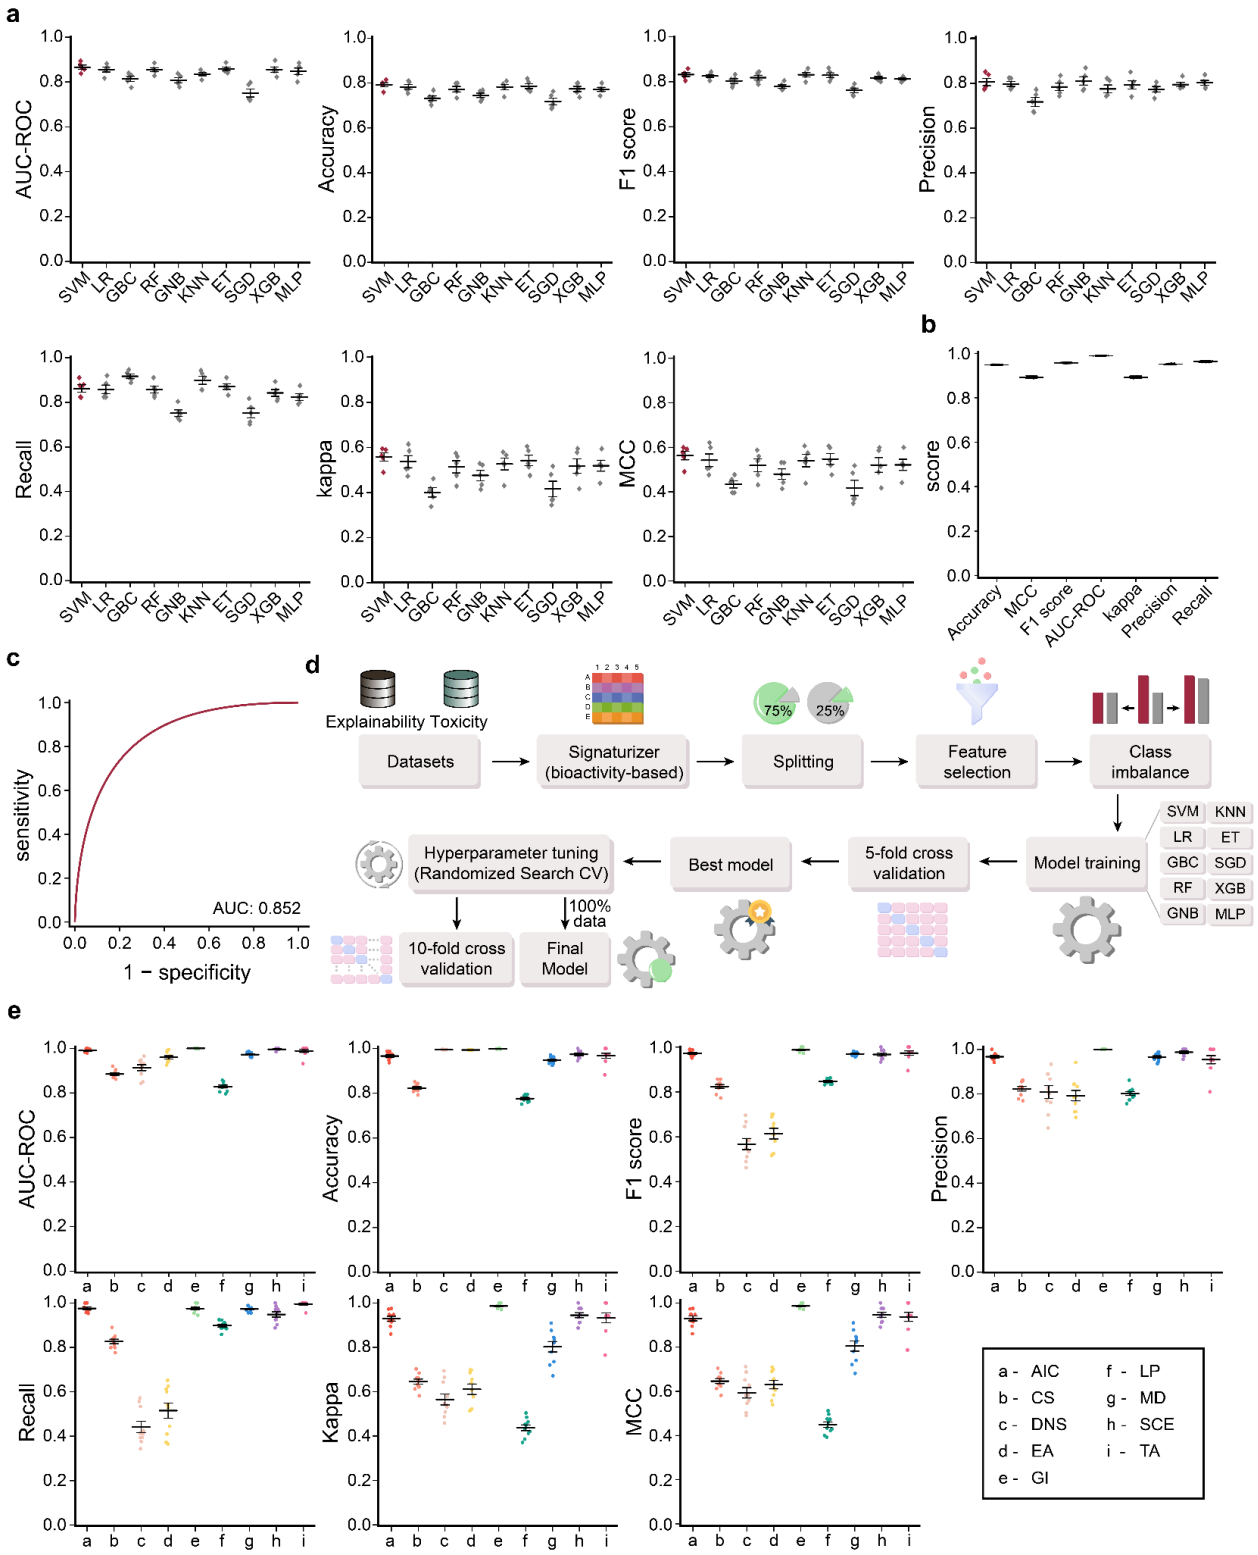

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

(a) Scatter–interval plots (n = 5 folds) summarizing five-fold cross-validation results across multiple performance

metrics, including AUC–ROC, accuracy, F1 score, precision, recall, Cohen's kappa, and Matthews correlation coefficient (MCC). Models were trained on experimentally validated geroprotectors and neutral compounds using the following classifiers: LR, SVM, GBC, RF, GNB, KNN, ET, SGD, XGB, and MLP. (b) Scatter–interval plot showing training performance metrics from leave-one-out cross-validation (LOOCV) of the SVM model (n = 972 folds). (c) AUC–ROC curve summarizing LOOCV testing performance of the SVM classifier. (d) Schematic overview of the model development pipeline for the Explainability and Toxicity modules. Signaturizer-derived bioactivity embeddings were used as molecular descriptors. Data were split into training (75%) and testing (25%) sets, Boruta was applied for feature selection, and SMOTE addressed class imbalance. Model selection was guided by five-fold cross-validation, followed by hyperparameter tuning and ten-fold cross-validation. Final models were trained on the full dataset. (e) Scatter–interval plots summarizing ten-fold cross-validation performance across multiple metrics. Whiskers represent ± SE, with bold lines indicating mean values.

Extended Data Figure 3: Model development and interpretability of Explainability and Toxicity modules

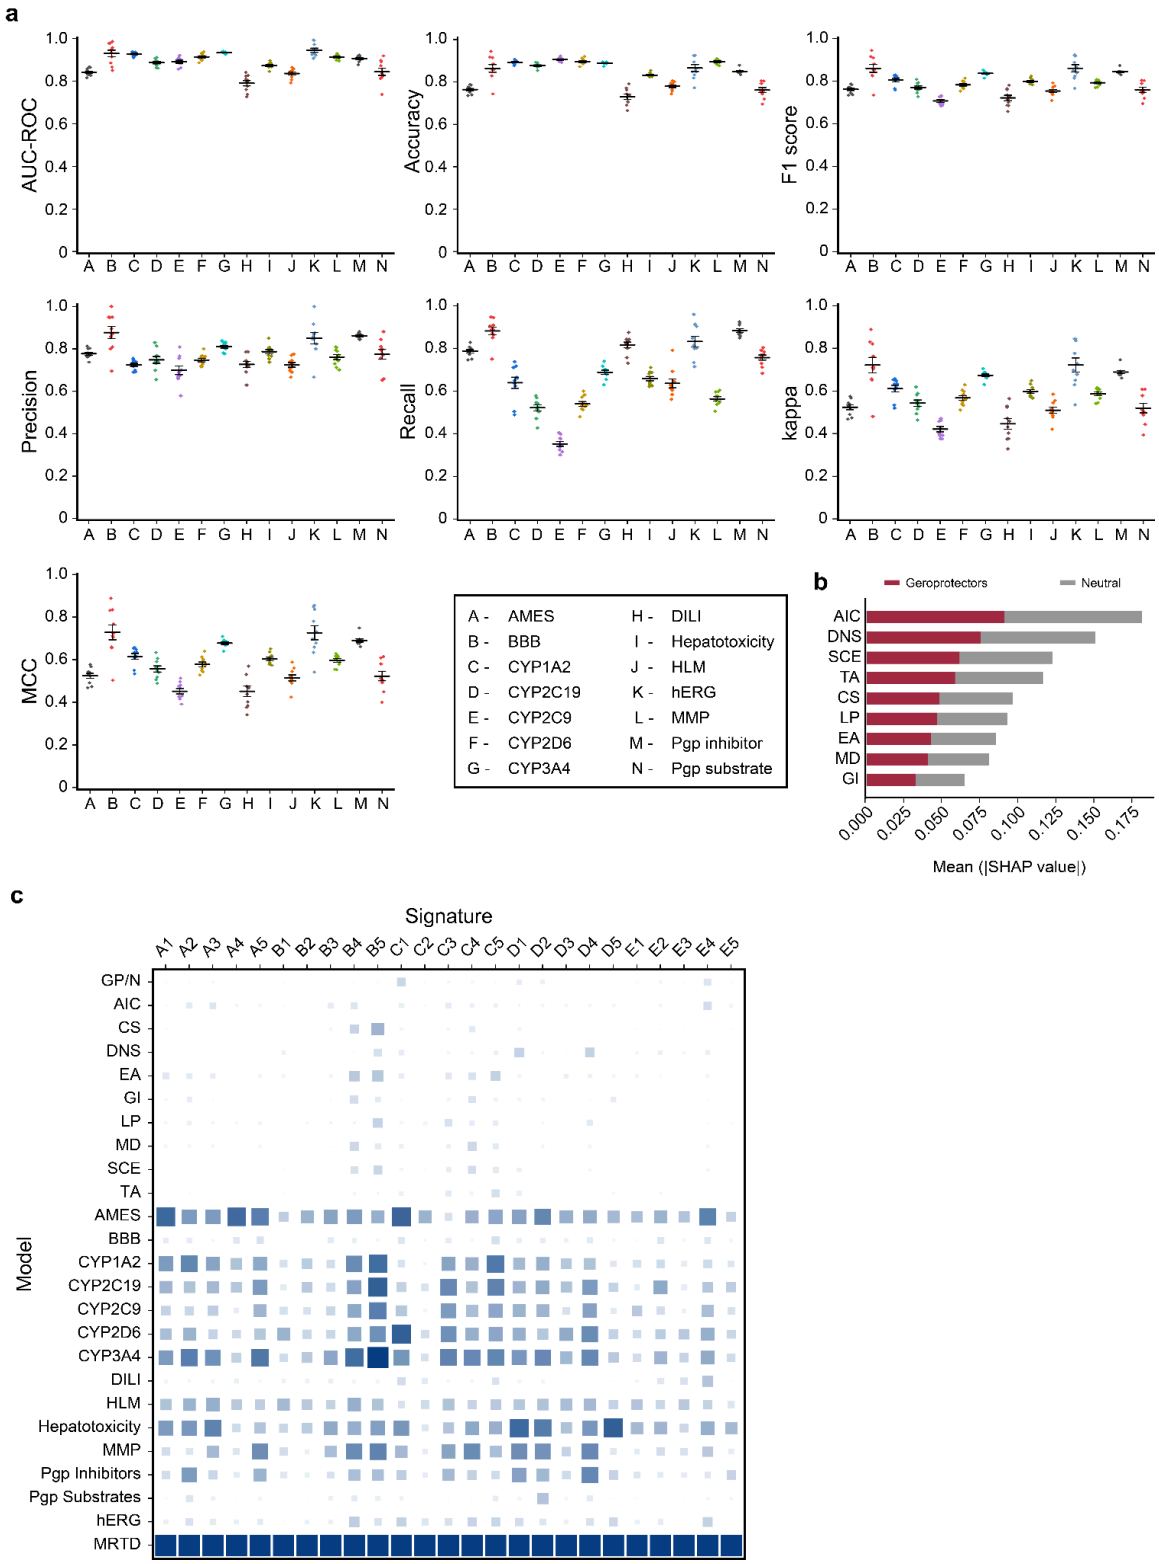

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

(a) Scatter–interval plots illustrating ten-fold cross-validation performance across standard classification metrics,

including AUC–ROC, accuracy, F1 score, precision, recall, Cohen's kappa, and MCC. All models utilized signaturizer-based bioactivity descriptors. (b) Stacked bar chart displaying the relative contribution of each aging-associated biological process to model predictions, quantified as mean absolute SHAP values. (c) Heatmap showing the proportion of selected signaturizer features retained after Boruta-based feature selection for each model.

Extended Data Figure 4: Large-scale identification of candidate geroprotectors

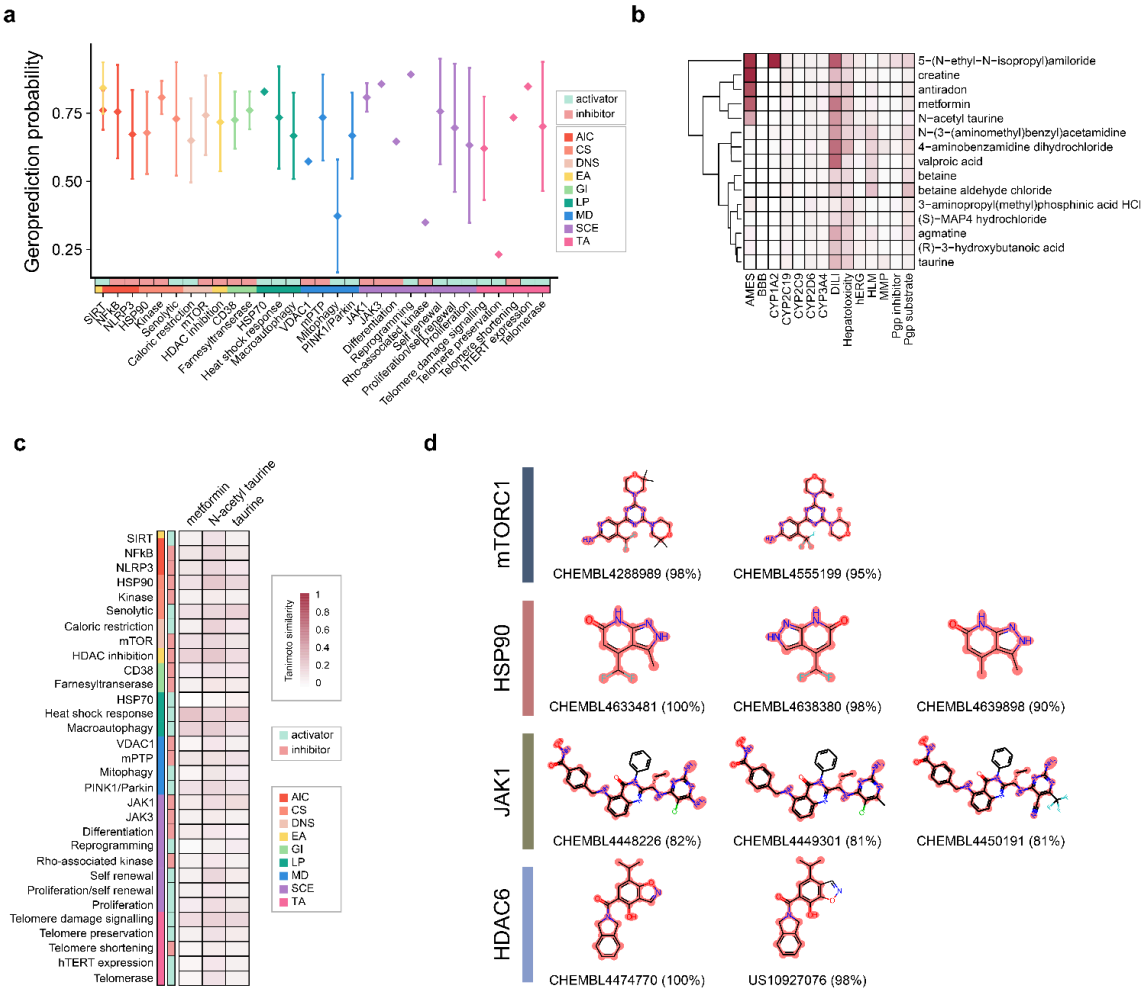

e

Database	Version	Initial counts	Post-processing counts	GP (≥ 0.904)
AfroDb	-	954	954	31.87 %
Aging Atlas	1.0	527	527	13.28 %
ChEMBL	ChEMBL31	1,918,718	1,914,746	2.69 %
ChemBridge	-	1,780	874	1.60 %
ChemDiv BBlocks (ChemDiv BB)	-	74,679	74,679	1.81 %
CMNPD	1.0	31,561	31,561	14.59 %
Dr. Duke's Phytochemicals (DDPEDB)	-	3,919	3,919	20.46 %
DSSTox	v2021	872,025	872,025	7.22 %
E. coli Metabolome (ECMDB)	2.0	3,757	3,757	26.91 %
FooDB	Pre-release 1.0	69,439	69,439	8.34 %
GDB	GDB13	977,468,301	975,902,199	1.68 %
gutMGene (GutMG)	Human	324	295	24.41 %
HMDB	5.0	217,308	217,308	40.39 %
IMPPAT 2.0 (IMPPAT)	2.0	17,968	13,468	28.99 %
MetaboAge	-	288	288	15.63 %
MolData	v1.1.0	1,429,988	1,429,988	0.88 %
PubChem	down-21/10/22	112,388,994	112,348,285	2.27 %
Drug Repurposing Hub (RepoHub)	down-24/03/20	6,717	6,717	5.51 %
YMDB	2.0	15,950	15,950	53.68 %
ZINC	ZINC12	45,195,702	45,194,425	1.05 %
Total		**1,139,718,899**	**1,138,101,404**	

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

(a) Mean–whisker plot summarizing average geroprediction probabilities for compounds predicted as negative

modulators of aging-associated processes (Explainability module, Class 1). (b) Heatmap depicting Toxicity module prediction scores for metformin, taurine, N-acetyl taurine, and their nearest neighbors in bioactivity embedding space. (c) Heatmap showing Tanimoto similarity between these compounds and known activators or inhibitors included in the Explainability module training set. (d) Schematic overview of Target module outputs, highlighting top ligand similarities for mTORC1, HSP90, JAK1, and HDAC6. (e) Table summarizing screened chemical databases, including database version, compound counts before and after preprocessing, and the number of high-confidence geroprotector predictions (probability ≥ 0.904).

Extended Data Figure 5: Mapping of AgeXtend-predicted geroprotective chemical space

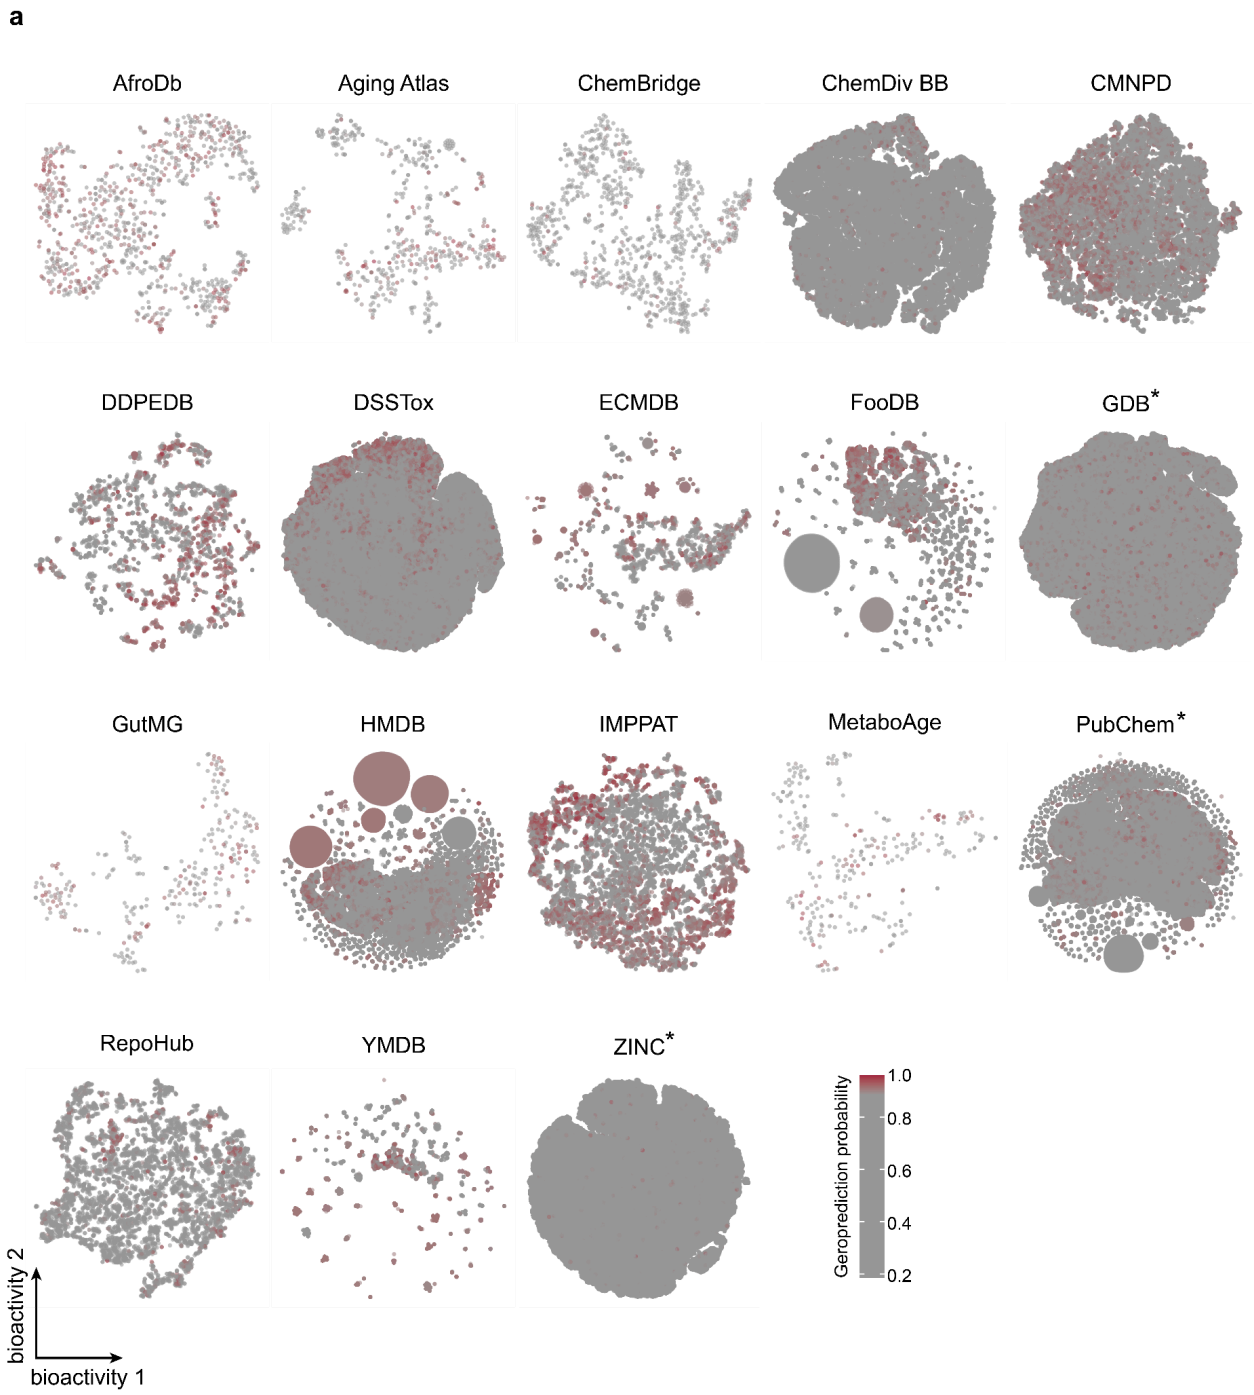

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

(a) t-SNE projections of high-confidence geroprotector predictions (probability ≥ 0.904) from individual databases, visualized in the latent space defined by Explainability module outputs. Databases marked with an asterisk (*) are represented by a subset of compounds.

Extended Data Figure 6: Mechanistic insights from the Explainability module

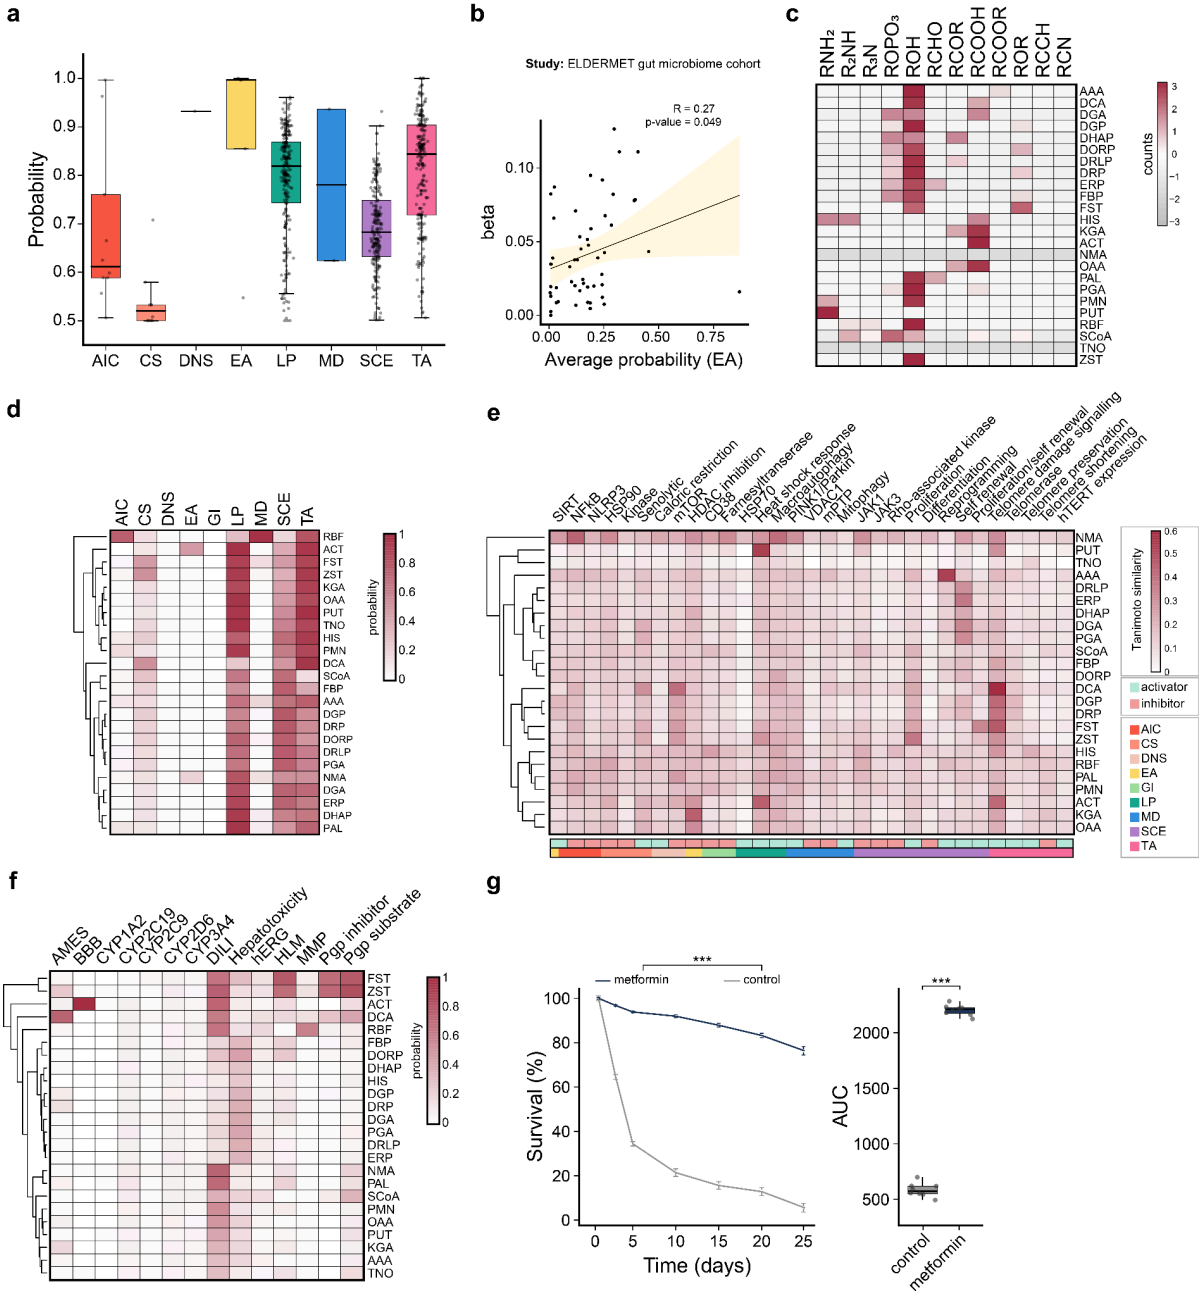

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

(a) Box plots showing probability distributions (>0.5) for nine aging-related biological processes across AgeXtend-predicted geroprotectors overlapping with differentially enriched or depleted metabolites from Aging Atlas. (b) Scatter plot illustrating the positive correlation between microbiome beta diversity and epigenetic alteration probabilities derived from the Explainability module in the ELDERMET cohort (Pearson's correlation, p = 0.049). (c) Heatmap representing relative enrichment or depletion of chemical functional groups across selected endogenous metabolites. (d) Heatmap of Explainability module prediction probabilities for all aging-associated processes across the indicated metabolites. (e) Heatmap of Tanimoto similarity between endogenous metabolites and

known pathway modulators from the training dataset. (f) Heatmap summarizing Toxicity module prediction probabilities for the same metabolites. (g) Yeast CLS survival analysis showing the geroprotective effect of metformin. Survival curves (left) represent mean ± SEM (n = 8 biological replicates), while box plots (right) show survival integrals computed as area under the survival curve. Statistical significance was assessed using a two-sided Mann–Whitney U test (p = 0.0002). For box plots in (a) and (g), center lines indicate medians, boxes denote interquartile ranges, and whiskers extend 1.5× IQR.

Extended Data Figure 7: Validation of geroprotective endogenous metabolites in yeast

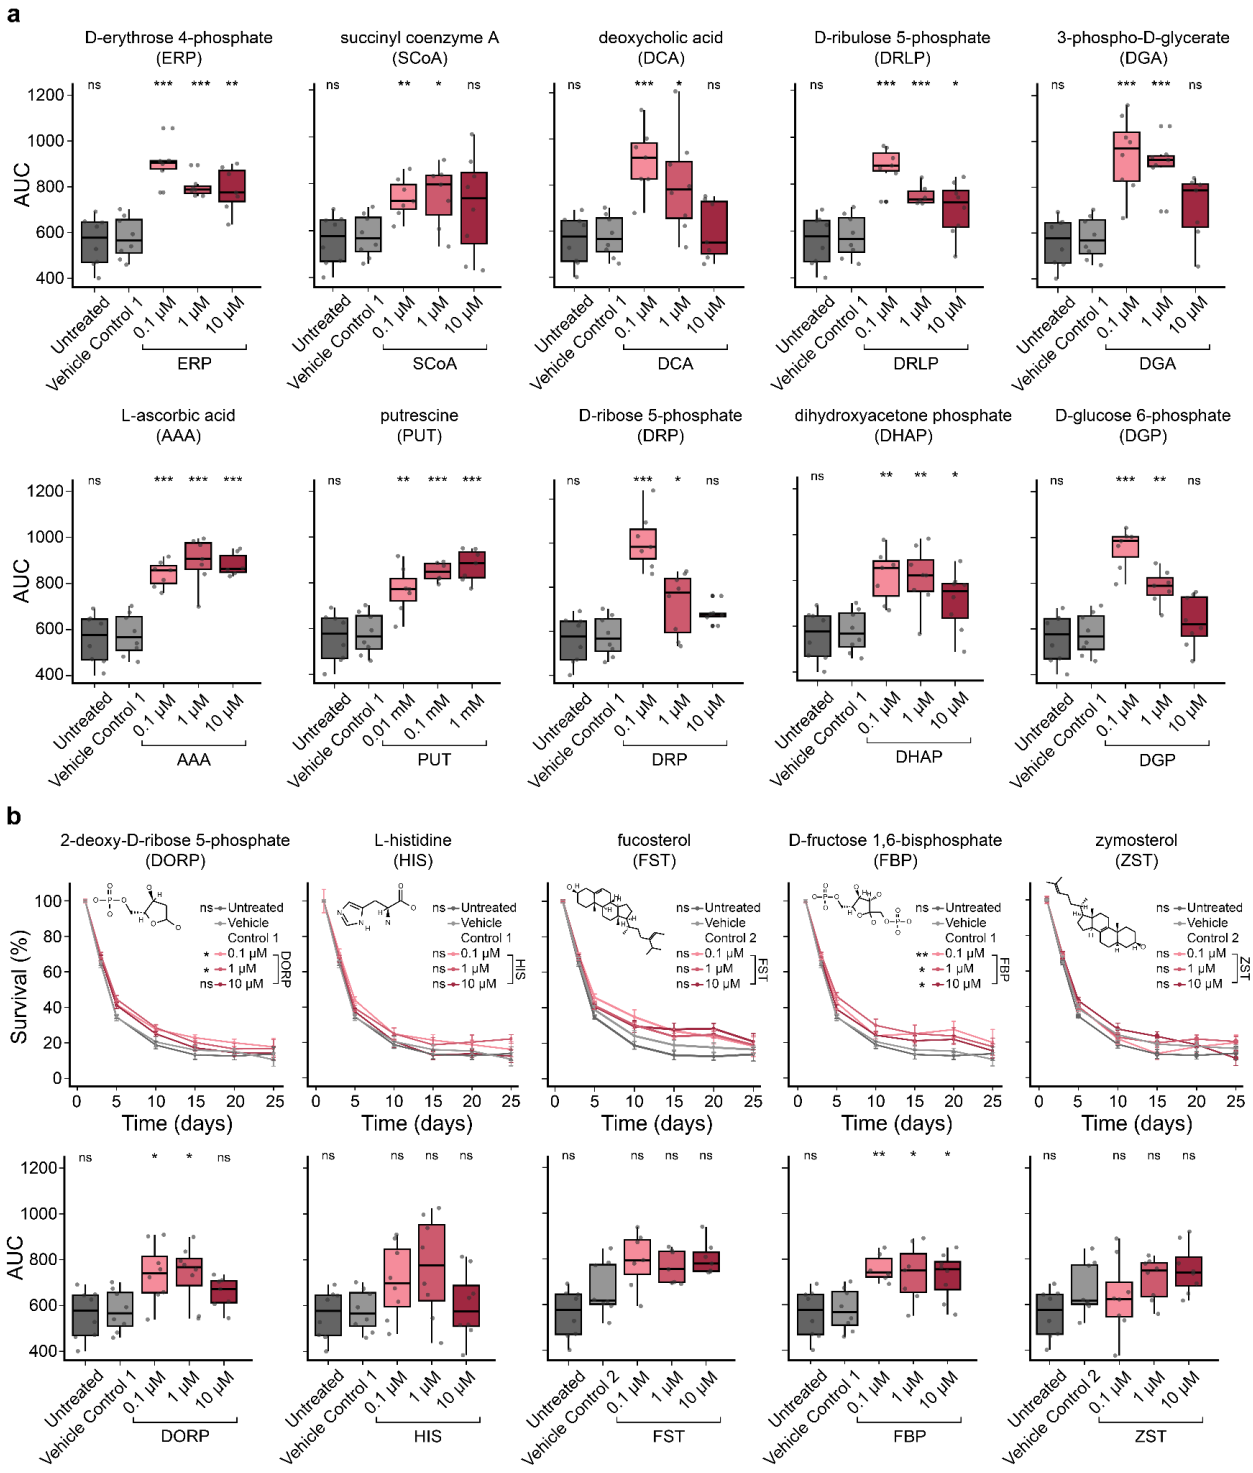

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

(a) Box plots showing survival integrals (area under CLS survival curves) across untreated, vehicle control, and metabolite-treated conditions at three concentrations (n = 8 biological replicates). Statistical significance was determined using a two-sided Mann–Whitney U test. (b) CLS survival curves demonstrating lifespan extension upon

treatment with AgeXtend-predicted metabolites. Error bars represent SEM across biological replicates. Corresponding box plots summarize survival integrals. Statistical comparisons were performed relative to vehicle controls using a two-sided Mann–Whitney U test.

Extended Data Figure 8: High validation success rate of AgeXtend predictions

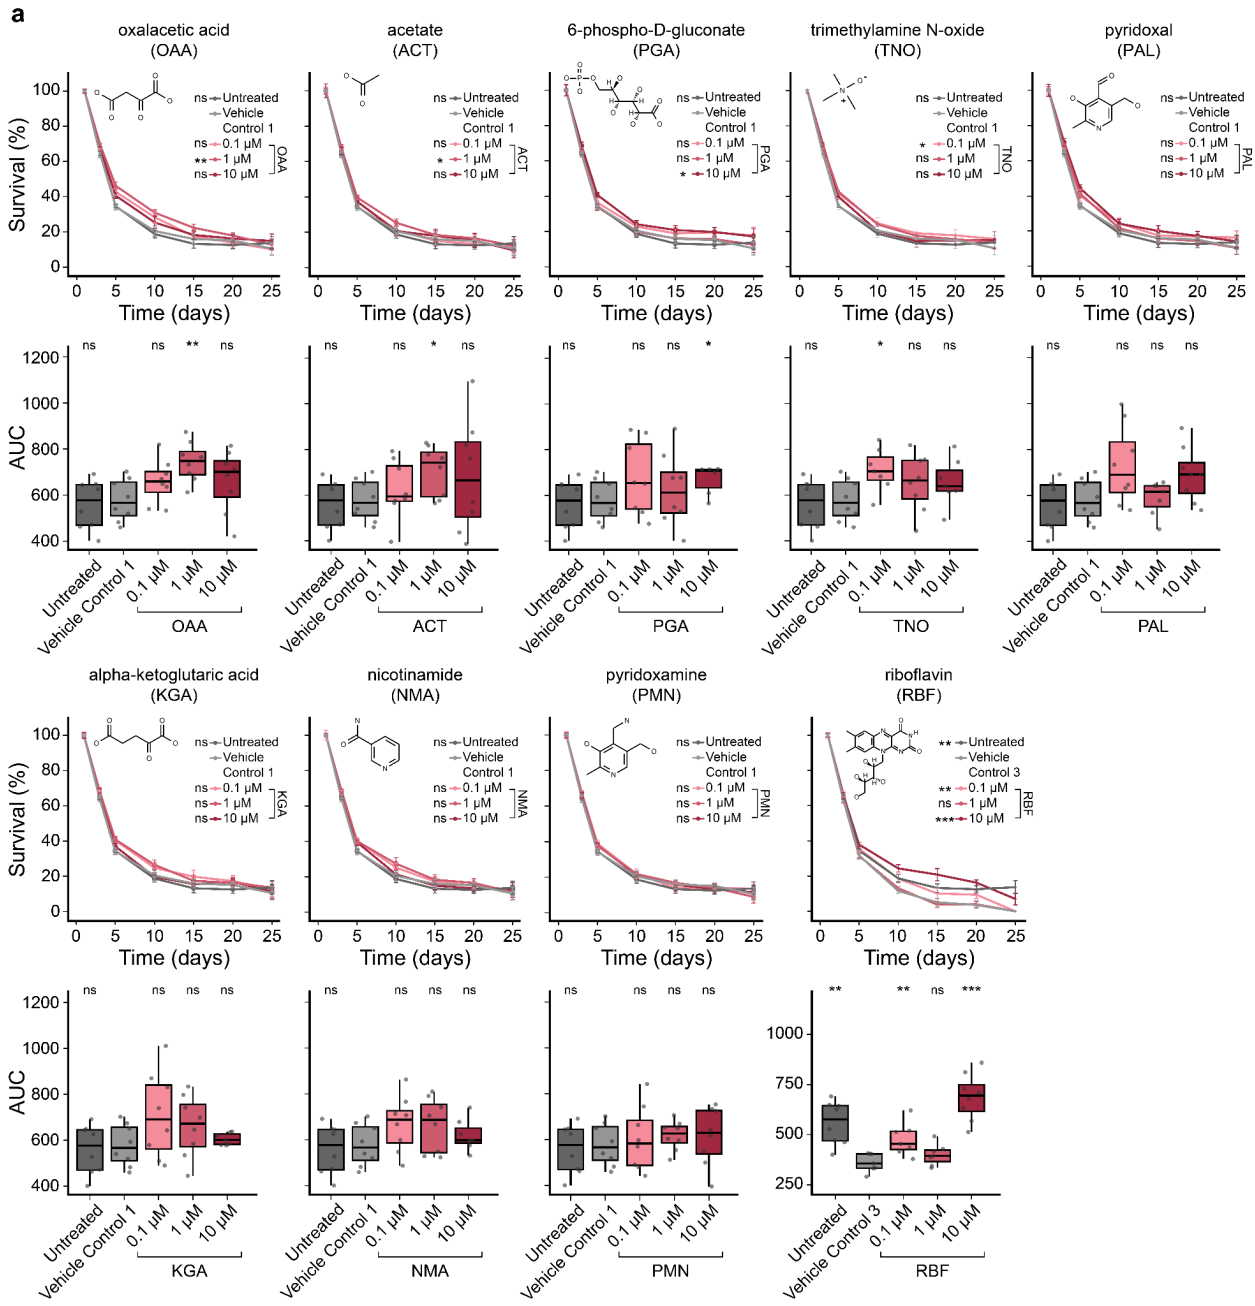

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

(a) Yeast CLS survival curves for AgeXtend-prioritized metabolites (n = 8 biological replicates). Mean survival ± SEM is shown, with box plots below indicating survival integrals. Statistical significance relative to vehicle controls was evaluated using a two-sided Mann–Whitney U test.

Extended Data Figure 9: Endogenous metabolites do not impair yeast growth

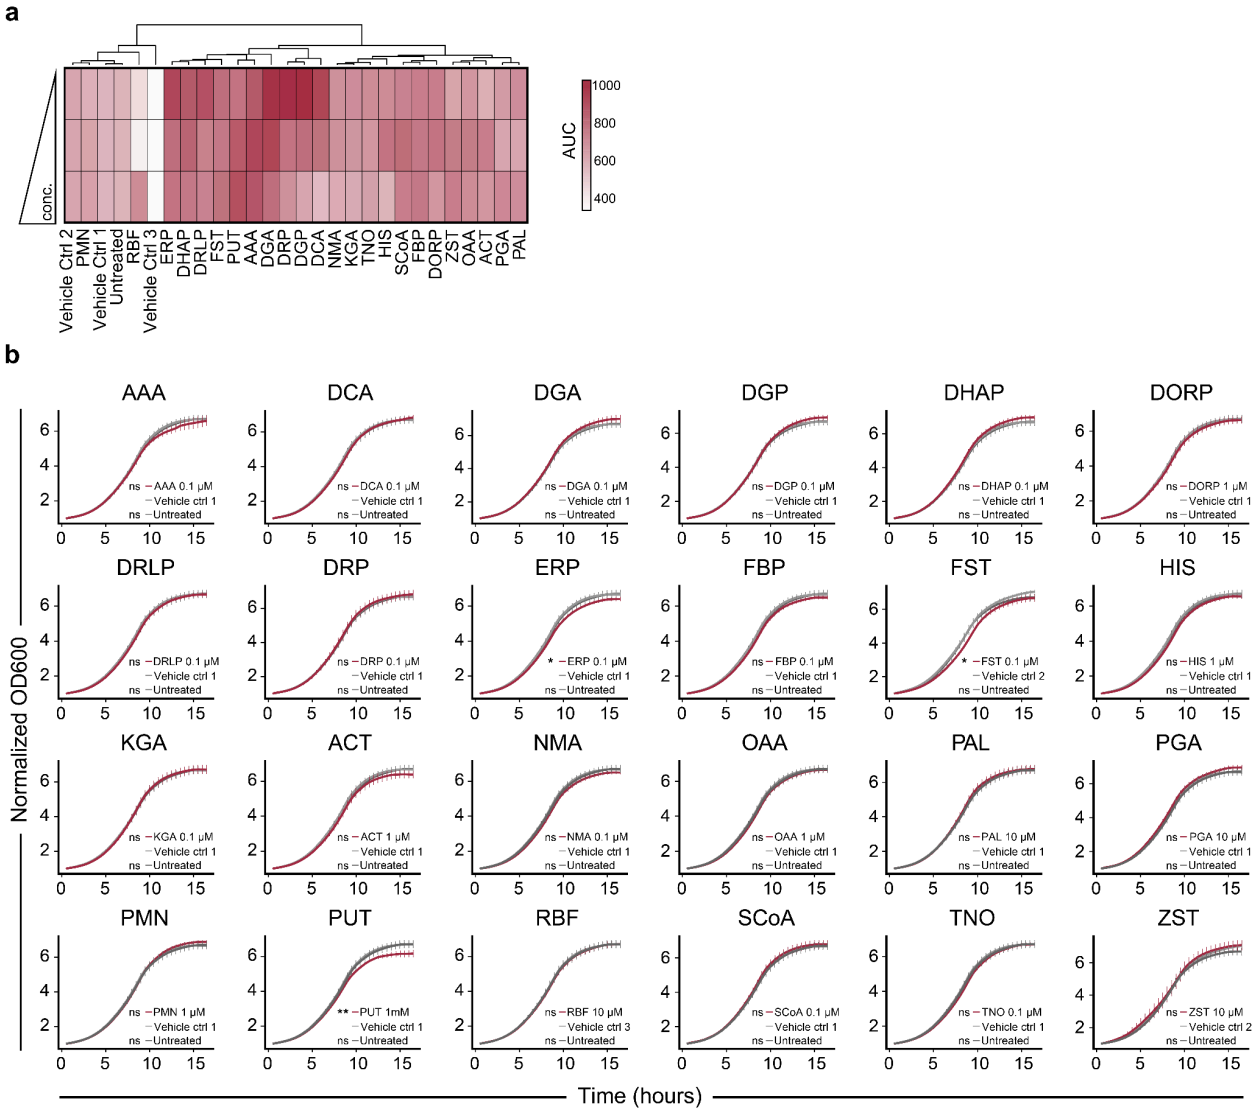

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

(a) Heatmap summarizing survival integrals from CLS assays across three tested concentrations of predicted metabolites. (b) Growth kinetics of yeast cultures treated with endogenous metabolites (n = 6 biological replicates), shown as normalized OD600 values ± SEM. Statistical significance of survival differences was assessed using a two-sided Mann–Whitney U test on area-under-curve values. Asterisks indicate $p < 0.05$; ns denotes non-significance.

Extended Data Figure 10: Toxicological profiling of endogenous metabolites

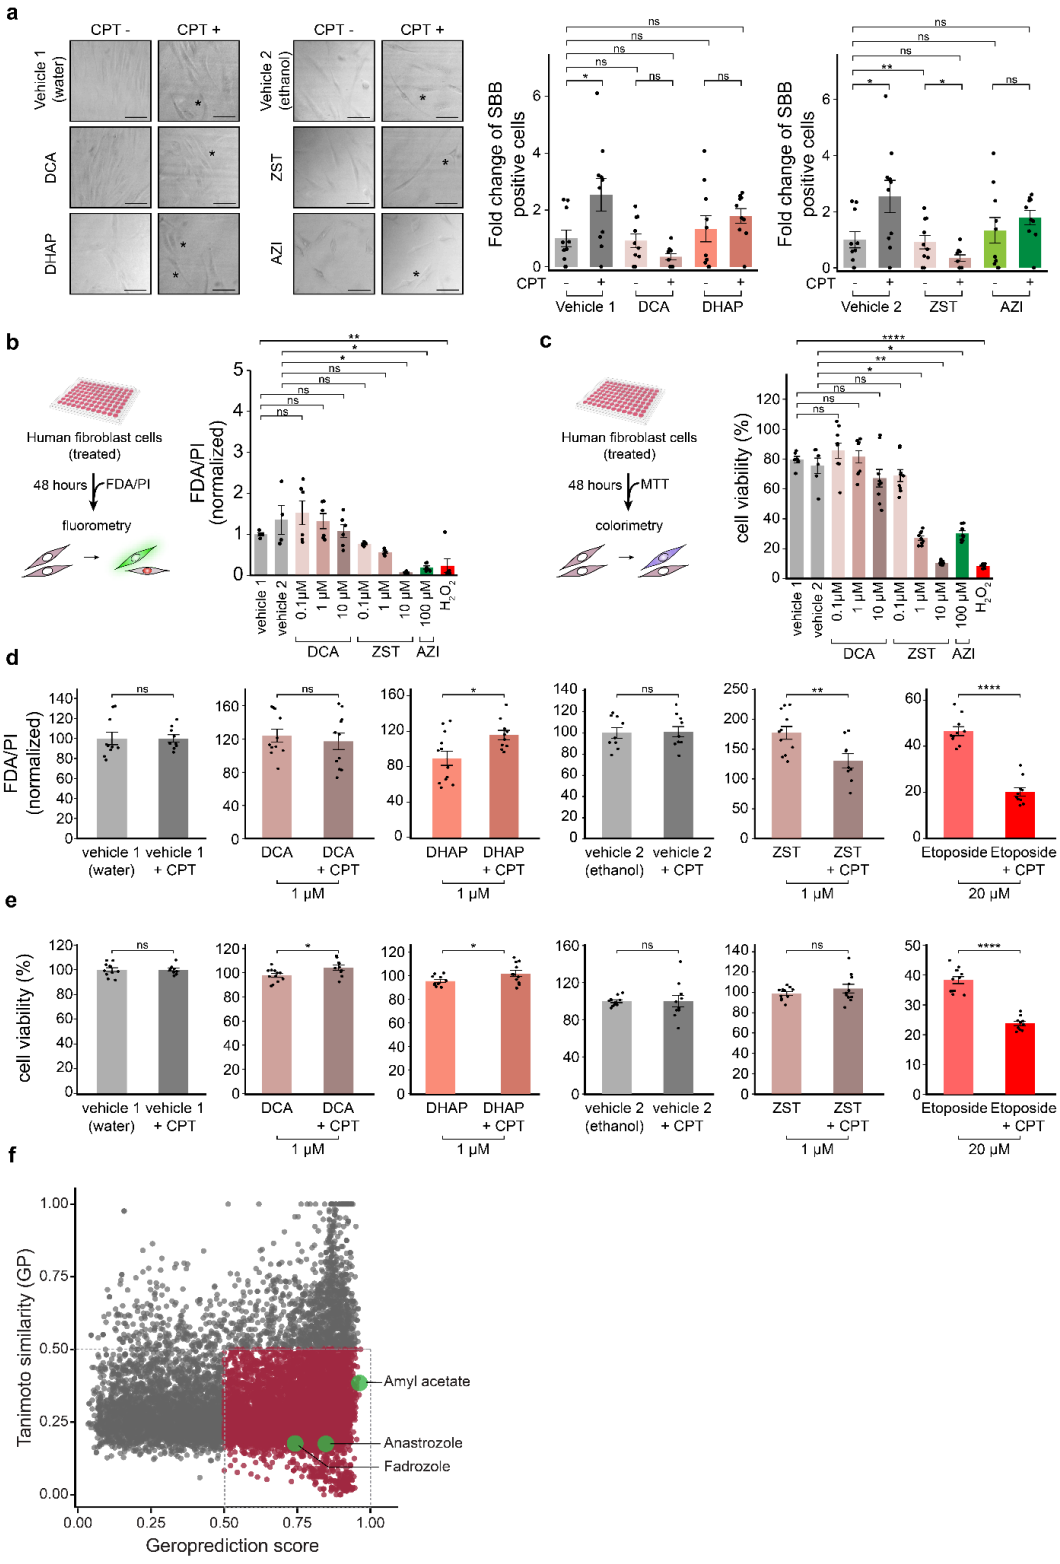

Reproduced from Arora S. 2025 et al., Nature Aging (permission granted)

(a) Representative micrographs of Sudan Black B staining in human fibroblasts under indicated conditions (scale

bar = 100 µm), with accompanying quantification of SBB-positive cells normalized to controls (n = 10 micrographs). (b) Experimental workflow and quantitative results for FDA/PI-based viability assessment. (c) Workflow and results for MTT-based viability measurements. (d–e) Quantification of FDA/PI and colorimetric assay outputs under indicated conditions. In panels (a–e), error bars represent SEM, and statistical significance was assessed using two-sided Welch's t-tests (95% confidence interval). (f) Scatter plot of DrugBank compounds positioned by AgeXtend geroprediction score (x-axis) and maximum Tanimoto similarity to known geroprotectors (y-axis). Compounds selected for validation of C. elegans lifespan are highlighted.

References

Abeliansky, Ana Lucia, Devin Erel, and Holger Strulik. 2020. “Aging in the USA: Similarities and Disparities across Time and Space.” *Scientific Reports* 10 (1): 14309.

“Advancing Geroscience Research – A Scoping Review of Regulatory Environments for Gerotherapeutics.” 2025. *The Journal of Nutrition, Health and Aging* 29 (9): 100637.

“Ageing and Health.” n.d. Accessed December 14, 2025. https://www.who.int/news-room/fact-sheets/detail/ageing-and-health.

Aging Atlas Consortium, Guang-Hui Liu, Yiming Bao, et al. 2020. “Aging Atlas: A Multi-Omics Database for Aging Biology.” *Nucleic Acids Research* 49 (D1): D825–D830.

“Aging: Therapeutics for a Healthy Future.” 2020. *Neuroscience & Biobehavioral Reviews* 108 (January): 453–458.

Ahmed, Mohamed. 2025. “Targeting Aging Pathways with Natural Compounds: A Review of Curcumin, Epigallocatechin Gallate, Thymoquinone, and Resveratrol.” *Immunity & Ageing* 22 (1): 1–18.

Amorim, João A., Giuseppe Coppotelli, Anabela P. Rolo, Carlos M. Palmeira, Jaime M. Ross, and David A. Sinclair. 2022. “Mitochondrial and Metabolic Dysfunction in Ageing and Age-Related Diseases.” *Nature Reviews Endocrinology* 18 (4): 243–258.

An, Yongpan, Qian Wang, Ke Gao, et al. 2025. “Epigenetic Regulation of Aging and Its Rejuvenation.” *MedComm* 6 (9): e70369.

Argentieri, M. Austin, Najaf Amin, Alejo J. Nevado-Holgado, et al. 2025. “Integrating the Environmental and Genetic Architectures of Aging and Mortality.” *Nature Medicine* 31 (3): 1016–1025.

Arora, Sakshi, Aayushi Mittal, Subhadeep Duari, et al. 2025. “Discovering Geroprotectors through the Explainable Artificial Intelligence-Based Platform AgeXtend.” *Nature Aging* 5 (1): 144–161.

“Artificial Intelligence (AI) in Drug Design and Discovery: A Comprehensive Review.” 2025. *In Silico Research in Biomedicine* 1 (January): 100049.

“Artificial Intelligence for Aging and Longevity Research: Recent Advances and Perspectives.” 2019. *Ageing Research Reviews* 49 (January): 49–66.

“Artificial Intelligence in Drug Discovery and Development.” 2021. *Drug Discovery Today* 26 (1): 80–93.

Axen, Seth D., Xi-Ping Huang, Elena L. Cáceres, Leo Gendelev, Bryan L. Roth, and Michael J. Keiser. 2017. “A Simple Representation of Three-Dimensional Molecular Structure.” *Journal of Medicinal Chemistry* 60 (17): 7393–7409.

Barardo, Diogo, Daniel Thornton, Harikrishnan Thoppil, et al. 2017. “The DrugAge Database of Aging-Related Drugs.” *Aging Cell* 16 (3): 594–597.

Barcena, Maria Luisa, Muhammad Aslam, Kristina Norman, Christiane Ott, and Yury Ladilov. 2025. “Role of AMPK and Sirtuins in Aging Heart: Basic and Translational Aspects.” *Aging and Disease* 16 (6): 3335–3360.

Barzilai, Nir, Jill P. Crandall, Stephen B. Kritchevsky, and Mark A. Espeland. 2016. “Metformin as a Tool to Target Aging.” *Cell Metabolism* 23 (6): 1060–1065.

Baur, Joseph A., Kevin J. Pearson, Nathan L. Price, et al. 2006. “Resveratrol Improves Health and Survival of Mice on a High-Calorie Diet.” *Nature* 444 (7117): 337–342.

Bell, Christopher G., Robert Lowe, Peter D. Adams, et al. 2019. “DNA Methylation Aging Clocks: Challenges and Recommendations.” *Genome Biology* 20 (1): 249.

Bernardes de Jesus, Bruno, Elsa Vera, Kerstin Schneeberger, et al. 2012. “Telomerase Gene Therapy in Adult and Old Mice Delays Aging and Increases Longevity without Increasing Cancer.” *EMBO Molecular Medicine* 4 (8): 691–704.

Bertoni, Martino, Miquel Duran-Frigola, Pau Badia-i-Mompel, et al. 2021. “Bioactivity Descriptors for Uncharacterized Chemical Compounds.” *Nature Communications* 12 (1): 3932.

Bibyk, Michael J., Melanie J. Campbell, and Amanda B. Hummon. 2021. “Mass Spectrometric Investigations of Caloric Restriction Mimetics.” *PROTEOMICS* 21 (9): 2000121.

“Biological Age across the Globe: 1990–2019.” 2025. *The Journal of the Economics of Ageing* 31 (June): 100573.

“Biological Age Predictors.” 2017. *EBioMedicine* 21 (July): 29–36.

Blagus, Rok, and Lara Lusa. 2013. “SMOTE for High-Dimensional Class-Imbalanced Data.” *BMC Bioinformatics* 14 (March): 106.

Bland, Jeffrey S. 2018. “Age as a Modifiable Risk Factor for Chronic Disease.” *Integrative Medicine (Encinitas, Calif.)* 17 (4): 16–19.

Cai, Hanxuan, Huimin Zhang, Duancheng Zhao, Jingxing Wu, and Ling Wang. 2022. “FP-GNN: A Versatile Deep Learning Architecture for Enhanced Molecular Property Prediction.” *Briefings in Bioinformatics* 23 (6). https://doi.org/10.1093/bib/bbac408.

“Caloric Restriction Mimetics against Age-Associated Disease: Targets, Mechanisms, and Therapeutic Potential.” 2019. *Cell Metabolism* 29 (3): 592–610.

“Calorie Restriction Mimetics: Can You Have Your Cake and Eat It, Too?” 2015. *Ageing Research Reviews* 20 (March): 46–62.

Campisi, Judith, Pankaj Kapahi, Gordon J. Lithgow, Simon Melov, John C. Newman, and Eric Verdin. 2019. “From Discoveries in Ageing Research to Therapeutics for Healthy Ageing.” *Nature* 571 (7764): 183–192.

Cantó, Carles, and Johan Auwerx. 2011. “Calorie Restriction: Is AMPK a Key Sensor and Effector?” *Physiology (Bethesda, Md.)* 26 (4): 214–224.

Cheng, Liang, Changlu Qi, Haixiu Yang, et al. 2022. “gutMGene: A Comprehensive Database for Target Genes of Gut Microbes and Microbial Metabolites.” *Nucleic Acids Research* 50 (D1): D795–D800.

Chen, Jianwei, Zude Liu, Hantao Wang, et al. 2021. “SIRT6 Enhances Telomerase Activity to Protect against DNA Damage and Senescence in Hypertrophic Ligamentum Flavum Cells from Lumbar Spinal Stenosis Patients.” *Aging* 13 (4): 6025–6040.

Chen, Ting, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. “A Simple Framework for Contrastive Learning of Visual Representations.” February 13. http://arxiv.org/abs/2002.05709.

Chung, Ki Wung, and Hae Young Chung. 2019. “The Effects of Calorie Restriction on Autophagy: Role on Aging Intervention.” *Nutrients* 11 (12): 2923.

Cohen, Alan A., Luigi Ferrucci, Tamàs Fülöp, et al. 2022. “A Complex Systems Approach to Aging Biology.” *Nature Aging* 2 (7): 580–591.

Conboy, Irina M., Michael J. Conboy, Amy J. Wagers, Eric R. Girma, Irving L. Weissman, and Thomas A. Rando. 2005. “Rejuvenation of Aged Progenitor Cells by Exposure to a Young Systemic Environment.” *Nature* 433 (7027): 760–764.

“Conserved Signaling Pathways Genetically Associated with Longevity across the Species.” 2019. *Biochimica et Biophysica Acta (BBA) - Molecular Basis of Disease* 1865 (7): 1745–1755.

Coppé, Jean-Philippe, Christopher K. Patil, Francis Rodier, et al. 2008. “Senescence-Associated Secretory Phenotypes Reveal Cell-Nonautonomous Functions of Oncogenic RAS and the p53 Tumor Suppressor.” *PLOS Biology* 6 (12): e301.

Dakik, Pamela, Mélissa McAuley, Marisa Chancharoen, et al. 2019. “Pairwise Combinations of Chemical Compounds That Delay Yeast Chronological Aging through Different Signaling Pathways Display Synergistic Effects on the Extent of Aging Delay.” *Oncotarget* 10 (3): 313–338.

De-Leon-Covarrubias, Ulises Edgardo, Jose Juan Perez-Trujillo, Sheila Adela Villa-Cedillo, et al. 2024. “Unlocking the Potential: Caloric Restriction, Caloric Restriction Mimetics, and Their Impact on Cancer Prevention and Treatment.” *Metabolites* 14 (8): 418.

Dhahbi, Joseph M., Patricia L. Mote, Gregory M. Fahy, and Stephen R. Spindler. 2005. “Identification of Potential Caloric Restriction Mimetics by Microarray Profiling.” *Physiological Genomics*, ahead of print, November 17. https://doi.org/10.1152/physiolgenomics.00069.2005.

“Dietary Polyphenols as Geroprotective Compounds: From Blue Zones to Hallmarks of Ageing.” 2025. *Ageing Research Reviews* 108 (June): 102733.

“Dietary Restriction in Rats and Mice: A Meta-Analysis and Review of the Evidence for Genotype-Dependent Effects on Lifespan.” 2012. *Ageing Research Reviews* 11 (2): 254–270.

Di Francesco, Andrea, Andrew G. Deighan, Lev Litichevskiy, et al. 2024. “Dietary Restriction Impacts Health and Lifespan of Genetically Diverse Mice.” *Nature* 634 (8034): 684–692.

DiLoreto, Race, and Coleen T. Murphy. 2017. “The Cell Biology of Aging.” *Molecular Biology of the Cell*, ahead of print, October 13. https://doi.org/10.1091/mbc.E14-06-1084.

Dolejs, Josef, and Petra Marešová. 2017. “Onset of Mortality Increase with Age and Age Trajectories of Mortality from All Diseases in the Four Nordic Countries.” *Clinical Interventions in Aging* 12 (January): 161–173.

Dönertaş, Handan Melike, Matías Fuentealba Valenzuela, Linda Partridge, and Janet M. Thornton. 2018. “Gene Expression-Based Drug Repurposing to Target Aging.” *Aging Cell* 17 (5): e12819.

Duval, Manuel X. 2018. “The Inadequacy of the Reductionist Approach in Discovering New Therapeutic Agents against Complex Diseases.” *Experimental Biology and Medicine (Maywood, N.J.)* 243 (12):

1004–1013.

Eggert, Ulrike S. 2013. “The Why and How of Phenotypic Small-Molecule Screens.” *Nature Chemical Biology* 9 (4): 206–209.

Eisenberg, Tobias, Mahmoud Abdellatif, Sabrina Schroeder, et al. 2016. “Cardioprotection and Lifespan Extension by the Natural Polyamine Spermidine.” *Nature Medicine* 22 (12): 1428–1438.

“Epigenetic Interplay between Histone Modifications and DNA Methylation in Gene Silencing.” 2008. *Mutation Research/Reviews in Mutation Research* 659 (1-2): 40–48.

Eslami Manoochehri, Hafez, Susmitha Sri Kadiyala, Javad Birjandtalab, and Mehrdad Nourani. 2018. “Feature Selection to Predict Compound’s Effect on Aging.” *Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, August 15, 419–427.

Evason, Kimberley, Cheng Huang, Idella Yamben, Douglas F. Covey, and Kerry Kornfeld. 2005. “Anticonvulsant Medications Extend Worm Life-Span.” *Science*, ahead of print, January 14. https://doi.org/10.1126/science.1105299.

Fang, Weicheng, Shuxin Chen, Xuejiao Jin, Shenkui Liu, Xiuling Cao, and Beidong Liu. 2023. “Metabolomics in Aging Research: Aging Markers from Organs.” *Frontiers in Cell and Developmental Biology* 11 (June): 1198794.

Ferreira, Fábio J. N., and Agnaldo S. Carneiro. 2025. “AI-Driven Drug Discovery: A Comprehensive Review.” *ACS Omega* 10 (23): 23889–23903.

Flanagan, Emily W., Jasper Most, Jacob T. Mey, and Leanne M. Redman. 2020. “Calorie Restriction and Aging in Humans.” *Annual Review of Nutrition* 40 (September): 105–133.

Fontana, Luigi, Linda Partridge, and Valter D. Longo. 2010. “Extending Healthy Life Span--from Yeast to Humans.” *Science (New York, N.Y.)* 328 (5976): 321–326.

Furrer, Regula, and Christoph Handschin. 2023. “Drugs, Clocks and Exercise in Ageing: Hype and Hope, Fact and Fiction.” *The Journal of Physiology* 601 (11): 2057–2068.

Gabandé-Rodríguez, Enrique, Manuel M. Gómez de Las Heras, and María Mittelbrunn. 2019. “Control of Inflammation by Calorie Restriction Mimetics: On the Crossroad of Autophagy and Mitochondria.” *Cells (Basel, Switzerland)* 9 (1): 82.

Garmany, Armin, and Andre Terzic. 2024. “Global Healthspan-Lifespan Gaps among 183 World Health Organization Member States.” *JAMA Network Open* 7 (12): e2450241.

Garmany, Armin, and Andre Terzic. 2025. “Healthspan-Lifespan Gap Differs in Magnitude and Disease Contribution across World Regions.” *Communications Medicine* 5 (1): 381.

Garmany, Armin, Satsuki Yamada, and Andre Terzic. 2021. “Longevity Leap: Mind the Healthspan Gap.” *Npj Regenerative Medicine* 6 (1): 57.

Gavrilov, Leonid A., and Natalia S. Gavrilova. 2019. “New Trend in Old-Age Mortality: Gompertzialization of Mortality Trajectory.” *Gerontology* 65 (5): 451–457.

“Genomic Instability as a Main Driving Factor of Unsuccessful Ageing: Potential for Translating the Use

of Micronuclei into Clinical Practice." 2021. *Mutation Research/Reviews in Mutation Research* 787 (January): 108359.

Ghosh, Tarini S., Mrinmoy Das, Ian B. Jeffery, and Paul W. O'Toole. 2020. "Adjusting for Age Improves Identification of Gut Microbiome Alterations in Multiple Diseases." *eLife* 9 (March). https://doi.org/10.7554/eLife.50240.

Ghosh, Tarini Shankar, Fergus Shanahan, and Paul W. O'Toole. 2022a. "The Gut Microbiome as a Modulator of Healthy Ageing." *Nature Reviews. Gastroenterology & Hepatology* 19 (9): 565–584.

Ghosh, Tarini Shankar, Fergus Shanahan, and Paul W. O'Toole. 2022b. "Toward an Improved Definition of a Healthy Microbiome for Healthy Aging." *Nature Aging* 2 (11): 1054–1069.

Gillespie, Zoe E., Joshua Pickering, and Christopher H. Eskiw. 2016. "Better Living through Chemistry: Caloric Restriction (CR) and CR Mimetics Alter Genome Function to Promote Increased Health and Lifespan." *Frontiers in Genetics* 7 (August): 142.

GitHub. n.d. "GitHub - Deepchem/deepchem: Democratizing Deep-Learning for Drug Discovery, Quantum Chemistry, Materials Science and Biology." Accessed December 16, 2025. https://github.com/deepchem/deepchem.

Gkioni, Lisonia, Tobias Nespital, Maarouf Baghdadi, et al. 2025. "The Geroprotectors Trametinib and Rapamycin Combine Additively to Extend Mouse Healthspan and Lifespan." *Nature Aging* 5 (7): 1249–1265.

Gompertz, Benjamin. 1825. "XXIV. On the Nature of the Function Expressive of the Law of Human Mortality, and on a New Mode of Determining the Value of Life Contingencies. In a Letter to Francis Baily, Esq. F. R. S. &c." *Philosophical Transactions of the Royal Society of London*, ahead of print, December 31. https://doi.org/10.1098/rstl.1825.0026.

Gonzalo, Susana. 2010. "Epigenetic Alterations in Aging." *Journal of Applied Physiology*, ahead of print, August 1. https://doi.org/10.1152/japplphysiol.00238.2010.

Gressler, A. Elisabeth, Houfu Leng, Heidi Zinecker, and Anna Katharina Simon. 2023. "Proteostasis in T Cell Aging." *Seminars in Immunology* 70 (101838): 101838.

Groff-Vindman, Cindy S., Benjamin D. Trump, Christopher L. Cummings, et al. 2025. "The Convergence of AI and Synthetic Biology: The Looming Deluge." *Npj Biomedical Innovations* 2 (1): 20.

Guan, Guangqi, Yaoxing Chen, and Yulan Dong. 2025. "Unraveling the AMPK-SIRT1-FOXO Pathway: The In-Depth Analysis and Breakthrough Prospects of Oxidative Stress-Induced Diseases." *Antioxidants (Basel, Switzerland)* 14 (1). https://doi.org/10.3390/antiox14010070.

Guo, Jun, Xiuqing Huang, Lin Dou, et al. 2022. "Aging and Aging-Related Diseases: From Molecular Mechanisms to Interventions and Treatments." *Signal Transduction and Targeted Therapy* 7 (1): 391.

Hakem, Razqallah. 2008. "DNA‐damage Repair; the Good, the Bad, and the Ugly." *The EMBO Journal* 27 (4): 589–605.

"Hallmarks of Health." 2021. *Cell* 184 (1): 33–63.

Harman, Ashley, and Tracy M. Bryan. 2024. "Telomere Maintenance and the DNA Damage Response: A Paradoxical Alliance." *Frontiers in Cell and Developmental Biology* 12 (October): 1472906.

Harrison, David E., Randy Strong, Zelton Dave Sharp, et al. 2009. “Rapamycin Fed Late in Life Extends Lifespan in Genetically Heterogeneous Mice.” *Nature* 460 (7253): 392–395.

Hegsted, D. M. 1990. “The Retardation of Aging and Disease by Dietary Restriction. Richard Weindruch , Roy L. Walford.” *The Quarterly Review of Biology* 65 (2): 256–256.

HelpAge International. 2024. “Population Ageing: Navigating the Demographic Shift.” July 11. https://www.helpage.org/news/population-ageing-navigating-the-demographic-shift/.

Herzog, Chiara M. S., Ludger J. E. Goeminne, Jesse R. Poganik, et al. 2024. “Challenges and Recommendations for the Translation of Biomarkers of Aging.” *Nature Aging* 4 (10): 1372–1383.

Hofer, Sebastian J., Sergio Davinelli, Martina Bergmann, Giovanni Scapagnini, and Frank Madeo. 2021. “Caloric Restriction Mimetics in Nutrition and Clinical Trials.” *Frontiers in Nutrition* 8 (September): 717343.

Ho, Kwok M., David J. Morgan, Mason Johnstone, and Cyrus Edibam. 2023. “Biological Age Is Superior to Chronological Age in Predicting Hospital Mortality of the Critically Ill.” *Internal and Emergency Medicine* 18 (7): 2019–2028.

Houtkooper, Riekelt H., Robert W. Williams, and Johan Auwerx. 2010. “Metabolic Networks of Longevity.” *Cell* 142 (1): 9–14.

Huang, Ruixue, and Ping-Kun Zhou. 2021. “DNA Damage Repair: Historical Perspectives, Mechanistic Pathways and Clinical Translation for Targeted Cancer Therapy.” *Signal Transduction and Targeted Therapy* 6 (1): 254.

Huang, Xinyi, Xingtong Pei, Weiyan Jian, and Mingming Xu. 2025. “Disability-Related Loss in Lifespan and Specific Social Determinants of Health Among Middle-Aged and Older Adults with Disabilities: Evidence from China’s Aging Population.” *Journal of Aging & Social Policy*, June 29, 1–19.

Hunt, Piper R., Tae Gen Son, Mark A. Wilson, et al. 2011. “Extension of Lifespan in C. Elegans by Naphthoquinones That Act through Stress Hormesis Mechanisms.” *PLOS ONE* 6 (7): e21922.

Ingram, Donald K., R. Michael Anson, Rafael de Cabo, et al. 2004. “Development of Calorie Restriction Mimetics as a Prolongevity Strategy.” *Annals of the New York Academy of Sciences* 1019 (June): 412–423.

Ingram, Donald K., Min Zhu, Jacek Mamczarz, et al. 2006. “Calorie Restriction Mimetics: An Emerging Research Field.” *Aging Cell* 5 (2): 97–108.

Irwin, John J., and Brian K. Shoichet. 2005. “ZINC--a Free Database of Commercially Available Compounds for Virtual Screening.” *Journal of Chemical Information and Modeling* 45 (1): 177–182.

Itahana, Koji, Goberdhan Dimri, and Judith Campisi. 2001. “Regulation of Cellular Senescence by p53.” *European Journal of Biochemistry* 268 (10): 2784–2791.

Jafari, Maryam, Adrián Macho-González, Antonio Diaz, et al. 2024. “Calorie Restriction and Calorie-Restriction Mimetics Activate Chaperone-Mediated Autophagy.” *Proceedings of the National Academy of Sciences of the United States of America* 121 (26): e2317945121.

Ji, Lixin, S. Michal Jazwinski, and Sangkyu Kim. 2021. “Frailty and Biological Age.” *Annals of Geriatric Medicine and Research* 25 (3): 141–149.

Ji, Shuaifei, Mingchen Xiong, Huating Chen, et al. 2023. “Cellular Rejuvenation: Molecular Mechanisms and Potential Therapeutic Interventions for Diseases.” *Signal Transduction and Targeted Therapy* 8 (1): 116.

Johnson, Simon C., Melana E. Yanos, Ernst-Bernhard Kayser, et al. 2013. “mTOR Inhibition Alleviates Mitochondrial Disease in a Mouse Model of Leigh Syndrome.” *Science (New York, N.Y.)* 342 (6165): 1524–1528.

Justice, Jamie, Jordan D. Miller, John C. Newman, et al. 2016. “Frameworks for Proof-of-Concept Clinical Trials of Interventions That Target Fundamental Aging Processes.” *The Journals of Gerontology. Series A, Biological Sciences and Medical Sciences* 71 (11): 1415–1423.

Justice, Jamie N., Anoop M. Nambiar, Tamar Tchkonia, et al. 2019. “Senolytics in Idiopathic Pulmonary Fibrosis: Results from a First-in-Human, Open-Label, Pilot Study.” *EBioMedicine* 40 (February): 554–563.

Kane, Alice E., and David A. Sinclair. 2019. “Epigenetic Changes during Aging and Their Reprogramming Potential.” *Critical Reviews in Biochemistry and Molecular Biology*, March 1, 61–83.

Kapsiani, Sofia, and Brendan J. Howlin. 2021. “Random Forest Classification for Predicting Lifespan-Extending Chemical Compounds.” *Scientific Reports* 11 (1): 13812.

Karagöz, Mustafa Fevzi, and A. Gülçin Sağdıçoğlu Celep. 2023. “The Effect of Caloric Restriction on Genetical Pathways.” *Food Science and Human Wellness* 12 (5): 1450–1457.

Kattuparambil, Aman Achuthan, Dheeraj Kumar Chaurasia, Shashank Shekhar, et al. 2025. “Exploring Chemical Space for ‘druglike’ Small Molecules in the Age of AI.” *Frontiers in Molecular Biosciences* 12 (March): 1553667.

Kennedy, Brian K., Shelley L. Berger, Anne Brunet, et al. 2014. “Geroscience: Linking Aging to Chronic Disease.” *Cell* 159 (4): 709–713.

Keshavarzi Arshadi, Arash, Milad Salem, Arash Firouzbakht, and Jiann Shiun Yuan. 2022. “MolData, a Molecular Benchmark for Disease and Target Based Machine Learning.” *Journal of Cheminformatics* 14 (1): 10.

Khosla, Prannay, Piotr Teterwak, Chen Wang, et al. 2020. “Supervised Contrastive Learning.” April 23. http://arxiv.org/abs/2004.11362.

Kornfeld, Kerry, and Kimberley Evason. 2006. “Effects of Anticonvulsant Drugs on Life Span.” *Archives of Neurology* 63 (4): 491–496.

Krenz, Astrid, and Holger Strulik. 2023. “Physiological Aging in India: The Role of the Epidemiological Transition.” *PLOS ONE* 18 (7): e0287259.

Kulkarni, Ameya S., Sandra Aleksic, David M. Berger, Felipe Sierra, George A. Kuchel, and Nir Barzilai. 2022. “Geroscience-Guided Repurposing of FDA-Approved Drugs to Target Aging: A Proposed Process and Prioritization.” *Aging Cell* 21 (4): e13596.

Kursa, Miron B., and Witold R. Rudnicki. 2010. “Feature Selection with the Boruta Package.” *Journal of Statistical Software* 36 (September): 1–13.

Lakshminarasimhan, Mahadevan, David Rauh, Mike Schutkowski, and Clemens Steegborn. 2013. “Sirt1 Activation by Resveratrol Is Substrate Sequence-Selective.” *Aging* 5 (3): 151–154.

Lamming, Dudley W., Lan Ye, David M. Sabatini, and Joseph A. Baur. 2013. “Rapalogs and mTOR Inhibitors as Anti-Aging Therapeutics.” *The Journal of Clinical Investigation* 123 (3): 980–989.

Lan, Jing, Hexiao Ding, Hongzhao Chen, et al. 2025. “Structure-Aware Contrastive Learning with Fine-Grained Binding Representations for Drug Discovery.” September 18. http://arxiv.org/abs/2509.14788.

Law, Vivian, Craig Knox, Yannick Djoumbou, et al. 2014. “DrugBank 4.0: Shedding New Light on Drug Metabolism.” *Nucleic Acids Research* 42 (Database issue): D1091–7.

Lee, Deborah J. W., Ajla Hodzic Kuerec, and Andrea B. Maier. 2024. “Targeting Ageing with Rapamycin and Its Derivatives in Humans: A Systematic Review.” *The Lancet. Healthy Longevity* 5 (2): e152–e162.

Lee, Shin-Hae, and Kyung-Jin Min. 2013. “Caloric Restriction and Its Mimetics.” *BMB Reports* 46 (4): 181–187.

Levine, Morgan E. 2012. “Modeling the Rate of Senescence: Can Estimated Biological Age Predict Mortality More Accurately Than Chronological Age?” *The Journals of Gerontology: Series A* 68 (6): 667–674.

Lin, Xiaokun, Lei Han, Jialei Weng, Kelai Wang, and Tongke Chen. 2018. “Rapamycin Inhibits Proliferation and Induces Autophagy in Human Neuroblastoma Cells.” *Bioscience Reports* 38 (6). https://doi.org/10.1042/BSR20181822.

Liu, James Y., Joshua Peeples, and Christie M. Sayes. 2024. “Evaluation of Machine Learning Based QSAR Models for the Classification of Lung Surfactant Inhibitors.” *Environment & Health*, ahead of print, September 20. https://doi.org/10.1021/envhealth.4c00118.

Liu, Tiqing, Yuhmei Lin, Xin Wen, Robert N. Jorissen, and Michael K. Gilson. 2007. “BindingDB: A Web-Accessible Database of Experimentally Determined Protein-Ligand Binding Affinities.” *Nucleic Acids Research* 35 (Database issue): D198–201.

Liu, Yu ’e, Chao Chen, Xinye Wang, et al. 2022. “An Epigenetic Role of Mitochondria in Cancer.” *Cells (Basel, Switzerland)* 11 (16): 2518.

Li, Xia, Chentao Li, Wanying Zhang, Yanan Wang, Pengxu Qian, and He Huang. 2023. “Inflammation and Aging: Signaling Pathways and Intervention Therapies.” *Signal Transduction and Targeted Therapy* 8 (1): 239.

Li, Zhe, Weiguang Zhang, Yuting Duan, et al. 2023. “Progress in Biological Age Research.” *Frontiers in Public Health* 11 (April): 1074274.

Longo, Valter D., and Mark P. Mattson. 2014. “Fasting: Molecular Mechanisms and Clinical Applications.” *Cell Metabolism* 19 (2): 181–192.

López-Otín, Carlos, Maria A. Blasco, Linda Partridge, Manuel Serrano, and Guido Kroemer. 2013. “The Hallmarks of Aging.” *Cell* 153 (6): 1194–1217.

López-Otín, Carlos, Maria A. Blasco, Linda Partridge, Manuel Serrano, and Guido Kroemer. 2023.

"Hallmarks of Aging: An Expanding Universe." *Cell* 186 (2): 243–278.

Lu, Jessica K., Weilan Wang, Muhammad Daniel Azlan Mahadzir, et al. 2025. "Digital Biomarkers of Ageing for Monitoring Physiological Systems in Community-Dwelling Adults." *The Lancet. Healthy Longevity* 6 (6): 100725.

Lushchak, Oleh, and Dmytro Gospodaryov. 2017. "CHAPTER 10. Mimetics of Caloric Restriction." In *Drug Discovery*. RSC Drug Discovery. Royal Society of Chemistry.

Madeo, Frank, Federico Pietrocola, Tobias Eisenberg, and Guido Kroemer. 2014. "Caloric Restriction Mimetics: Towards a Molecular Definition." *Nature Reviews Drug Discovery* 13 (10): 727–740.

Ma, Lina, Rong Wang, Hongjuan Wang, Yaxin Zhang, and Zhiwei Zhao. 2020. "Long-Term Caloric Restriction Activates the Myocardial SIRT1/AMPK/PGC-1α Pathway in C57BL/6J Male Mice." *Food & Nutrition Research* 64 (0). https://doi.org/10.29219/fnr.v64.3668.

Mannick, Joan B., and Dudley W. Lamming. 2023. "Targeting the Biology of Aging with mTOR Inhibitors." *Nature Aging* 3 (6): 642–660.

Mariño, Guillermo, Federico Pietrocola, Frank Madeo, and Guido Kroemer. 2014. "Caloric Restriction Mimetics: Natural/physiological Pharmacological Autophagy Inducers." *Autophagy* 10 (11): 1879–1882.

Marino, Nicola, Guido Putignano, Simone Cappilli, et al. 2023. "Towards AI-Driven Longevity Research: An Overview." *Frontiers in Aging* 4 (March): 1057204.

Martinez-Lopez, Nuria, Diana Athonvarangkul, and Rajat Singh. 2015. "Autophagy and Aging." *Longevity Genes*, 73–87.

Masoro, E. J. 2000. "Caloric Restriction and Aging: An Update." *Experimental Gerontology* 35 (3): 299–305.

"Mechanisms Linking Mitochondrial Dysfunction and Proteostasis Failure." 2020. *Trends in Cell Biology* 30 (4): 317–328.

Mercken, Evi M., Miriam Capri, Bethany A. Carboneau, et al. 2017. "Conserved and Species-Specific Molecular Denominators in Mammalian Skeletal Muscle Aging." *Npj Aging and Mechanisms of Disease* 3 (1): 8.

"Mitochondrial-Epigenetic Crosstalk in Environmental Toxicology." 2017. *Toxicology* 391 (November): 5–17.

Moqri, Mahdi, Chiara Herzog, Jesse R. Poganik, et al. 2024. "Validation of Biomarkers of Aging." *Nature Medicine* 30 (2): 360–372.

Moriwaki, Hirotomo, Yu-Shi Tian, Norihito Kawashita, and Tatsuya Takagi. 2018. "Mordred: A Molecular Descriptor Calculator." *Journal of Cheminformatics* 10 (February): 4.

Moskalev, Alexey, Elizaveta Chernyagina, Anna Kudryavtseva, and Mikhail Shaposhnikov. 2017. "Geroprotectors: A Unified Concept and Screening Approaches." *Aging and Disease* 8 (3): 354–363.

Moskalev, Alexey, Elizaveta Chernyagina, Vasily Tsvetkov, et al. 2016. "Developing Criteria for Evaluation of Geroprotectors as a Key Stage toward Translation to the Clinic." *Aging Cell* 15 (3):

407–415.

Most, Jasper, Valeria Tosti, Leanne M. Redman, and Luigi Fontana. 2017. “Calorie Restriction in Humans: An Update.” *Ageing Research Reviews* 39 (October): 36–45.

Mouchiroud, Laurent, Laurent Molin, Nicolas Dallière, and Florence Solari. 2010. “Life Span Extension by Resveratrol, Rapamycin, and Metformin: The Promise of Dietary Restriction Mimetics for an Healthy Aging.” *BioFactors (Oxford, England)* 36 (5): 377–382.

Musa, Aliyu, Laleh Soltan Ghoraie, Shu-Dong Zhang, et al. 2018. “A Review of Connectivity Map and Computational Approaches in Pharmacogenomics.” *Briefings in Bioinformatics* 19 (3): 506–523.

Mutz, Julian, Raquel Iniesta, and Cathryn M. Lewis. 2024. “Metabolomic Age (MileAge) Predicts Health and Life Span: A Comparison of Multiple Machine Learning Algorithms.” *Science Advances*, ahead of print, December 20. https://doi.org/10.1126/sciadv.adp3743.

“NAD+ in Aging: Molecular Mechanisms and Translational Implications.” 2017. *Trends in Molecular Medicine* 23 (10): 899–916.

Neuberg, Leland Gerson. 2003. “CAUSALITY: MODELS, REASONING, AND INFERENCE, by Judea Pearl, Cambridge University Press, 2000.” *Econometric Theory* 19 (04). https://doi.org/10.1017/s0266466603004109.

Niazi, Sarfaraz K. 2025. “Artificial Intelligence in Small-Molecule Drug Discovery: A Critical Review of Methods, Applications, and Real-World Outcomes.” *Pharmaceuticals (Basel, Switzerland)* 18 (9): 1271.

Nousis, Lambros, Panagiotis Kanavaros, and Alexandra Barbouti. 2023. “Oxidative Stress-Induced Cellular Senescence: Is Labile Iron the Connecting Link?” *Antioxidants (Basel, Switzerland)* 12 (6). https://doi.org/10.3390/antiox12061250.

Ntie-Kang, Fidele, Denis Zofou, Smith B. Babiaka, et al. 2013. “AfroDb: A Select Highly Potent and Diverse Natural Product Library from African Medicinal Plants.” *PLOS ONE* 8 (10): e78085.

Oh, Juhyun, Yang David Lee, and Amy J. Wagers. 2014. “Stem Cell Aging: Mechanisms, Regulators and Therapeutic Opportunities.” *Nature Medicine* 20 (8): 870–880.

Pálovics, Róbert, Andreas Keller, Nicholas Schaum, et al. 2022. “Molecular Hallmarks of Heterochronic Parabiosis at Single-Cell Resolution.” *Nature* 603 (7900): 309–314.

Pan, Yan, Hongxia Cai, Fang Ye, et al. 2025. “ElixirSeeker: A Machine Learning Framework Utilizing Fusion Molecular Fingerprints for the Discovery of Lifespan-Extending Compounds.” *Aging Cell* 24 (8): e70116.

Parkhitko, Andrey A., Elizabeth Filine, and Marc Tatar. 2023. “Combinatorial Interventions in Aging.” *Nature Aging* 3 (10): 1187–1200.

Partridge, Linda. 2010. “The New Biology of Ageing.” *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 365 (1537): 147–154.

Pearson, Kevin J., Joseph A. Baur, Kaitlyn N. Lewis, et al. 2008. “Resveratrol Delays Age-Related Deterioration and Mimics Transcriptional Aspects of Dietary Restriction without Extending Life Span.” *Cell Metabolism* 8 (2): 157–168.

Peoples, Jessica N., Anita Saraf, Nasab Ghazal, Tyler T. Pham, and Jennifer Q. Kwong. 2019. “Mitochondrial Dysfunction and Oxidative Stress in Heart Disease.” *Experimental & Molecular Medicine* 51 (12): 1–13.

Perri, Giorgia, Chloe French, César Agostinis-Sobrinho, et al. 2025. “An Expert Consensus Statement on Biomarkers of Aging for Use in Intervention Studies.” *The Journals of Gerontology. Series A, Biological Sciences and Medical Sciences* 80 (5). https://doi.org/10.1093/gerona/glae297.

“Perspectives of the Potential Implications of Polyphenols in Influencing the Interrelationship between Oxi-Inflammatory Stress, Cellular Senescence and Immunosenescence during Aging.” 2020. *Trends in Food Science & Technology* 98 (April): 41–52.

Rao, Mohan, Eric McDuffie, and Clifford Sachs. 2023. “Artificial Intelligence/machine Learning-Driven Small Molecule Repurposing via off-Target Prediction and Transcriptomics.” *Toxics* 11 (10): 875.

Ravussin, Eric, Leanne M. Redman, James Rochon, et al. 2015. “A 2-Year Randomized Controlled Trial of Human Caloric Restriction: Feasibility and Effects on Predictors of Health Span and Longevity.” *The Journals of Gerontology. Series A, Biological Sciences and Medical Sciences* 70 (9): 1097–1104.

“Recent Advances in the Field of Caloric Restriction Mimetics and Anti-Aging Molecules.” 2021. *Ageing Research Reviews* 66 (March): 101240.

Reddy, P. H., and T. P. Reddy. 2011. “Mitochondria as a Therapeutic Target for Aging and Neurodegenerative Diseases.” *Current Alzheimer Research* 8 (4): 393–409.

Redman, Leanne M., and Eric Ravussin. 2011. “Caloric Restriction in Humans: Impact on Physiological, Psychological, and Behavioral Outcomes.” *Antioxidants & Redox Signaling* 14 (2): 275–287.

Ribeiro, Caio, Christopher K. Farmer, João Pedro de Magalhães, and Alex A. Freitas. 2022. “Predicting Lifespan-Extending Chemical Compounds with Machine Learning and Biologically Interpretable Features.” In *bioRxiv*. November 22. https://doi.org/10.1101/2022.11.20.517230.

Rivero-Segura, Nadia Alejandra, Emmanuel Alejandro Zepeda-Arzate, Selma Karime Castillo-Vazquez, et al. 2024. “Exploring the Geroprotective Potential of Nutraceuticals.” *Nutrients* 16 (17). https://doi.org/10.3390/nu16172835.

Rolland, Yves, Felipe Sierra, Luigi Ferrucci, et al. 2023. “Challenges in Developing Geroscience Trials.” *Nature Communications* 14 (1): 5038.

Rose, Karen C., Laura N. Gitlin, and Marie P. Dennis. 2010. “Readiness to Use Compensatory Strategies among Older Adults with Functional Difficulties.” *International Psychogeriatrics* 22 (8): 1225–1239.

Roth, G. S., D. K. Ingram, and M. A. Lane. 2001. “Caloric Restriction in Primates and Relevance to Humans.” *Annals of the New York Academy of Sciences* 928 (1): 305–315.

ŞAHiN, Emrullah, Naciye Nur Arslan, and Durmuş Özdemir. 2024. “Unlocking the Black Box: An in-Depth Review on Interpretability, Explainability, and Reliability in Deep Learning.” *Neural Computing and Applications* 37 (2): 859–965.

Sanada, Fumihiro, Shinichiro Hayashi, and Ryuichi Morishita. 2025. “Targeting the Hallmarks of Aging: Mechanisms and Therapeutic Opportunities.” *Frontiers in Cardiovascular Medicine* 12 (July): 1631578.

Santiago-de-la-Cruz, Jose Alberto, Nadia Alejandra Rivero-Segura, and Juan Carlos Gomez-Verjan. 2025. "Structure-Based Machine Learning Screening Identifies Natural Product Candidates as Potential Geroprotectors." *Journal of Cheminformatics* 17 (1): 106.

Saretzki, Gabriele. 2025. "The Role of Telomeres in Senescence, Aging and Disease: Fiction and Reality." *OBM Geriatrics* 9 (3): 1–52.

Schyman, Patric, Ruifeng Liu, Valmik Desai, and Anders Wallqvist. 2017. "vNN Web Server for ADMET Predictions." *Frontiers in Pharmacology* 8 (December): 889.

Sciarretta, Sebastiano, Maurizio Forte, Francesca Castoldi, et al. 2021. "Caloric Restriction Mimetics for the Treatment of Cardiovascular Diseases." *Cardiovascular Research* 117 (6): 1434–1449.

"Senolytics Decrease Senescent Cells in Humans: Preliminary Report from a Clinical Trial of Dasatinib plus Quercetin in Individuals with Diabetic Kidney Disease." 2019. *EBioMedicine* 47 (September): 446–456.

Seo, Yonghak, Samuel Kingsley, Griffin Walker, Michelle A. Mondoux, and Heidi A. Tissenbaum. 2018. "Metabolic Shift from Glycogen to Trehalose Promotes Lifespan and Healthspan in Caenorhabditis Elegans." *Proceedings of the National Academy of Sciences of the United States of America* 115 (12): E2791–E2800.

Shen, Yao, Nairui Fan, Shu-Xia Ma, Xin Cheng, Xuesong Yang, and Guang Wang. 2025. "Gut Microbiota Dysbiosis: Pathogenesis, Diseases, Prevention, and Therapy." *MedComm* 6 (5): e70168.

Shintani, Hideya, Tomoya Shintani, Hisashi Ashida, and Masashi Sato. 2018. "Calorie Restriction Mimetics: Upstream-Type Compounds for Modulating Glucose Metabolism." *Nutrients* 10 (12): 1821.

Singh, Parminder, Kishore Gollapalli, Stefano Mangiola, et al. 2023. "Taurine Deficiency as a Driver of Aging." *Science*, ahead of print, June 9. https://doi.org/10.1126/science.abn9257.

Singh, Sandeep Kumar. 2025. "The Potential Mechanism and the Role of Antioxidants in Mitigating Oxidative Stress in Alzheimer's Disease." *Frontiers in Bioscience-Landmark* 30 (2). https://doi.org/10.31083/FBL25551.

Smith, Daniel L., Tim R. Nagy, and David B. Allison. 2010. "Calorie Restriction: What Recent Results Suggest for the Future of Ageing Research." *European Journal of Clinical Investigation* 40 (5): 440–450.

Soliman, Ghada A. 2022. "Intermittent Fasting and Time-Restricted Eating Role in Dietary Interventions and Precision Nutrition." *Frontiers in Public Health* 10 (October): 1017254.

Stephenne, X., M. Foretz, N. Taleux, et al. 2011. "Metformin Activates AMP-Activated Protein Kinase in Primary Human Hepatocytes by Decreasing Cellular Energy Status." *Diabetologia* 54 (12): 3101–3110.

Stolbov, Leonid, Anastasia Rudik, Alexey Lagunin, Dmitry Druzhilovskiy, Dmitry Filimonov, and Vladimir Poroikov. 2025. "In Silico Assessment of Potential Geroprotectors: From Separate Endpoints to Complex Pharmacotherapeutic Effects." *International Journal of Molecular Sciences* 26 (18). https://doi.org/10.3390/ijms26188858.

Szőke, Kitti, Beáta Bódi, Zoltán Hendrik, et al. 2023. "Rapamycin Treatment Increases Survival,

Autophagy Biomarkers and Expression of the Anti-Aging Klotho Protein in Elderly Mice." *Pharmacology Research & Perspectives* 11 (3): e01091.

Taormina, Giusi, and Mario G. Mirisola. 2014. "Calorie Restriction in Mammals and Simple Model Organisms." *BioMed Research International* 2014 (1): 308690.

Tao, Xue, Ziman Zhu, Liguo Wang, et al. 2024. "Biomarkers of Aging and Relevant Evaluation Techniques: A Comprehensive Review." *Aging and Disease* 15 (3): 977–1005.

Taylor, Rebecca C., and Andrew Dillin. 2011. "Aging as an Event of Proteostasis Collapse." *Cold Spring Harbor Perspectives in Biology* 3 (5): a004440.

Templeman, Nicole M., and Coleen T. Murphy. 2018. "Regulation of Reproduction and Longevity by Nutrient-Sensing Pathways." *The Journal of Cell Biology* 217 (1): 93–106.

Tenchov, Rumiana, Janet M. Sasso, Xinmei Wang, and Qiongqiong Angela Zhou. 2023. "Aging Hallmarks and Progression and Age-Related Diseases: A Landscape View of Research Advancement." *ACS Chemical Neuroscience*, ahead of print, December 14. https://doi.org/10.1021/acschemneuro.3c00531.

Tenchov, Rumiana, Janet M. Sasso, Xinmei Wang, and Qiongqiong Angela Zhou. 2024. "Antiaging Strategies and Remedies: A Landscape of Research Progress and Promise." *ACS Chemical Neuroscience*, ahead of print, January 12. https://doi.org/10.1021/acschemneuro.3c00532.

"The Aging Proteostasis Decline: From Nematode to Human." 2021. *Experimental Cell Research* 399 (2): 112474.

"The Health-Adjusted Dependency Ratio as a New Global Measure of the Burden of Ageing: A Population-Based Study." 2022. *The Lancet Healthy Longevity* 3 (5): e332–e338.

Thomas M. Cover, and Joy A. Thomas. 2006. *Elements of Information Theory: Cover/elements of Information Theory, Second Edition*. 2nd ed. John Wiley & Sons. PDF.

Tomar, Rachana, and Luke Erber. 2023. "Endogenous Metabolites and Genome Instability in Aging and Disease." *Chemical Research in Toxicology* 36 (12): 1830–1833.

Trisal, Anchal, and Abhishek Kumar Singh. 2024. "Clinical Insights on Caloric Restriction Mimetics for Mitigating Brain Aging and Related Neurodegeneration." *Cellular and Molecular Neurobiology* 44 (1): 1–47.

Vega Magdaleno, Gustavo Daniel, and Joao Pedro de Magalhaes. 2024. "Pleiotropy and Disease Interactors: The Dual Nature of Genes Linking Ageing and Ageing-Related Diseases." In *bioRxiv*. May 22. https://doi.org/10.1101/2024.05.20.594975.

Wang, Jie, Fanglin Shao, Qing Xin Yu, et al. 2025. "The Common Hallmarks and Interconnected Pathways of Aging, Circadian Rhythms, and Cancer: Implications for Therapeutic Strategies." *Research*, ahead of print, March 5. https://doi.org/10.34133/research.0612.

Wang, Lifen, Jason Karpac, Heinrich Jasper, Shireen A. Davies, Julian A. T. Dow, and Ken Lukowiak. 2014. "Promoting Longevity by Maintaining Metabolic and Proliferative Homeostasis." *Journal of Experimental Biology* 217 (1): 109–118.

Wang, Tongzhou, and Phillip Isola. 2020. "Understanding Contrastive Representation Learning through

Alignment and Uniformity on the Hypersphere." May 20. http://arxiv.org/abs/2005.10242.

Wang, Yu-Jen, Francesco Paneni, Sokrates Stein, and Christian M. Matter. 2021. "Modulating Sirtuin Biology and Nicotinamide Adenine Diphosphate Metabolism in Cardiovascular Disease—From Bench to Bedside." *Frontiers in Physiology* 12 (October): 755060.

Warren, Alison, Yvonne Nyavor, Nikkia Zarabian, Aidan Mahoney, and Leigh A. Frame. 2024. "The Microbiota-Gut-Brain-Immune Interface in the Pathogenesis of Neuroinflammatory Diseases: A Narrative Review of the Emerging Literature." *Frontiers in Immunology* 15 (May): 1365673.

Wassermann, Anne Mai, Eugen Lounkine, Dominic Hoepfner, et al. 2015. "Dark Chemical Matter as a Promising Starting Point for Drug Lead Discovery." *Nature Chemical Biology* 11 (12): 958–966.

"Website." n.d.-a. https://doi.org/10.1016/S0140-6736(15)00516-4.

"Website." n.d.-b. https://doi.org/10.1016/S0140-6736(15)00516-4.

"Website." n.d.-c. https://doi.org/10.1136/bmjgh-2024-015173.

"Website." n.d.-d. https://doi.org/10.1093/gerona/glad034.

"Website." n.d.-e. https://doi.org/10.1093/lifemedi/lnac018.

"Website." n.d.-f. https://doi.org/10.1016/j.tem.2022.01.007.

"Website." n.d.-g. https://doi.org/10.3390/ijms26199643.

Weinert, Brian T., and Poala S. Timiras. 2003. "Invited Review: Theories of Aging." *Journal of Applied Physiology*, ahead of print, October 1. https://doi.org/10.1152/japplphysiol.00288.2003.

World Health Organization. 2015. *World Report on Ageing and Health*. World Health Organization.

Yang, Minghao, Zehua Wang, Zizhuo Yan, Wenxiang Wang, Qian Zhu, and Changlong Jin. 2024. "DNASimCLR: A Contrastive Learning-Based Deep Learning Approach for Gene Sequence Data Classification." *BMC Bioinformatics* 25 (1): 328.

Yang, Yi, Sicong Chen, Yuqing Zhang, et al. 2017. "Induction of Autophagy by Spermidine Is Neuroprotective via Inhibition of Caspase 3-Mediated Beclin 1 Cleavage." *Cell Death & Disease* 8 (4): e2738–e2738.

Yao, Rufan, Zhenhua Shen, Xinyi Xu, et al. 2024. "Knowledge Mapping of Graph Neural Networks for Drug Discovery: A Bibliometric and Visualized Analysis." *Frontiers in Pharmacology* 15 (May): 1393415.

Yarmolenko, Mariia, and Brendan Howlin. 2022. *Extreme Gradient Boosting Algorithm Classification for Predicting Lifespan-Extending Chemical Compounds*. October 25. https://doi.org/10.21203/rs.3.rs-2199002/v1.

Yessenkyzy, Assylzhan, Timur Saliev, Marina Zhanaliyeva, et al. 2020. "Polyphenols as Caloric-Restriction Mimetics and Autophagy Inducers in Aging Research." *Nutrients* 12 (5): 1344.

Ye, Xiaolan, James M. Linton, Nicholas J. Schork, Linda B. Buck, and Michael Petrascheck. 2014. "A Pharmacological Network for Lifespan Extension in Caenorhabditis Elegans." *Aging Cell* 13 (2):

206–215.

Zeng, Xiangxiang, Siyi Zhu, Yuan Hou, et al. 2020. “Network-Based Prediction of Drug–target Interactions Using an Arbitrary-Order Proximity Embedded Deep Forest.” *Bioinformatics* 36 (9): 2805–2812.

Zhang, Kexin, Yujie Ma, Youhong Luo, et al. 2023. “Metabolic Diseases and Healthy Aging: Identifying Environmental and Behavioral Risk Factors and Promoting Public Health.” *Frontiers in Public Health* 11 (October): 1253506.

Zhang, Yingwei, Shikai Yan, Xiang Gao, et al. 2012. “Metabonomic Investigation on the Protective Effects of Rosiglitazone and Caloric Restriction for Renal Senescence in a Rat Model.” *Aging Clinical and Experimental Research* 24 (5): 430–438.

Zhao, Yang, Matthew Simon, Andrei Seluanov, and Vera Gorbunova. 2023. “DNA Damage and Repair in Age-Related Inflammation.” *Nature Reviews. Immunology* 23 (2): 75–89.

Zheng, Zikai, Jiaming Li, Tianzi Liu, et al. 2024. “DNA Methylation Clocks for Estimating Biological Age in Chinese Cohorts.” *Protein & Cell* 15 (8): 575–593.

Zhou, Guangfeng, Domnita-Valeria Rusnac, Hahnbeom Park, et al. 2024. “An Artificial Intelligence Accelerated Virtual Screening Platform for Drug Discovery.” *Nature Communications* 15 (1): 7761.

Zinecker, Heidi, and Anna Katharina Simon. 2022. “Autophagy Takes It All – Autophagy Inducers Target Immune Aging.” *Disease Models & Mechanisms* 15 (1): dmm049345.